



US007660800B2

(12) **United States Patent**
Prahlad et al.

(10) **Patent No.:** **US 7,660,800 B2**
(45) **Date of Patent:** **Feb. 9, 2010**

(54) **SYSTEMS AND METHODS FOR
CLASSIFYING AND TRANSFERRING
INFORMATION IN A STORAGE NETWORK**

(75) Inventors: **Anand Prahlad**, East Brunswick, NJ
(US); **Jeremy A. Schwartz**, Red Bank,
NJ (US); **David Ngo**, Shrewsbury, NJ
(US); **Brian Brockway**, Shrewsbury, NJ
(US); **Marcus S. Muller**, Tinton Falls,
NJ (US)

5,005,122 A	4/1991	Griffin et al.
5,093,912 A	3/1992	Dong et al.
5,133,065 A	7/1992	Cheffetz et al.
5,193,154 A	3/1993	Kitajima et al.
5,212,772 A	5/1993	Masters
5,226,157 A	7/1993	Nakano et al.
5,239,647 A	8/1993	Anglin et al.

(Continued)

(73) Assignee: **CommVault Systems, Inc.**, Oceanport,
NJ (US)

FOREIGN PATENT DOCUMENTS

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 144 days.

EP 0259912 3/1988

(Continued)

(21) Appl. No.: **11/564,153**

OTHER PUBLICATIONS

(22) Filed: **Nov. 28, 2006**

(65) Prior Publication Data

US 2007/0185916 A1 Aug. 9, 2007

Arneson, David A., "Development of Omniserver," Control Data
Corporation, Tenth IEEE Symposium on Mass Storage Systems,
May 1990, 'Crisis in Mass Storage' Digest of Papers, pp. 88-93,
Monterey, CA.

(Continued)

Related U.S. Application Data

(60) Provisional application No. 60/740,686, filed on Nov.
28, 2005, provisional application No. 60/752,203,
filed on Dec. 19, 2005.

Primary Examiner—James Trujillo
Assistant Examiner—Bruce M. Moser
(74) *Attorney, Agent, or Firm*—Perkins Coie LLP

(51) Int. Cl.

G06F 7/00 (2006.01)
G06F 17/30 (2006.01)
G06F 12/00 (2006.01)

(52) **U.S. Cl.** **707/10; 707/200; 709/201**

(58) **Field of Classification Search** **707/10,**
707/200; 709/201

See application file for complete search history.

(56) References Cited

U.S. PATENT DOCUMENTS


4,686,620 A 8/1987 Ng
4,995,035 A 2/1991 Cole et al.

(57) ABSTRACT

Systems and methods for data classification to facilitate and
improve data management within an enterprise are described.
The disclosed systems and methods evaluate and define data
management operations based on data characteristics rather
than data location, among other things. Also provided are
methods for generating a data structure of metadata that
describes system data and storage operations. This data struc-
ture may be consulted to determine changes in system data
rather than scanning the data files themselves.

13 Claims, 22 Drawing Sheets

100



US 7,660,800 B2

Page 2

U.S. PATENT DOCUMENTS

5,241,668 A	8/1993	Eastridge et al.	6,421,711 B1	7/2002	Blumenau et al.
5,241,670 A	8/1993	Eastridge et al.	6,421,779 B1	7/2002	Kuroda et al.
5,276,860 A	1/1994	Fortier et al.	6,430,575 B1	8/2002	Dourish et al.
5,276,867 A	1/1994	Kenley et al.	6,438,586 B1	8/2002	Hass et al.
5,287,500 A	2/1994	Stoppani, Jr.	6,487,561 B1	11/2002	Ofek et al.
5,321,816 A	6/1994	Rogan et al.	6,487,644 B1	11/2002	Huebsch et al.
5,333,315 A	7/1994	Saether et al.	6,519,679 B2	2/2003	Devireddy et al.
5,347,653 A	9/1994	Flynn et al.	6,538,669 B1	3/2003	Lagueux, Jr. et al.
5,410,700 A	4/1995	Fecteau et al.	6,542,909 B1	4/2003	Tamer et al.
5,448,724 A	9/1995	Hayashi et al.	6,542,972 B2	4/2003	Ignatius et al.
5,491,810 A	2/1996	Allen	6,564,228 B1	5/2003	O'Connor
5,495,607 A	2/1996	Pisello et al.	6,581,143 B2	6/2003	Gagne et al.
5,504,873 A	4/1996	Martin et al.	6,658,436 B2	12/2003	Oshinsky et al.
5,519,865 A	5/1996	Kondo et al.	6,658,526 B2	12/2003	Nguyen et al.
5,544,345 A	8/1996	Carpenter et al.	6,732,124 B1	5/2004	Koseki et al.
5,544,347 A	8/1996	Yanai et al.	6,772,164 B2	8/2004	Reinhardt
5,559,957 A	9/1996	Balk	6,775,790 B2	8/2004	Reuter et al.
5,619,644 A	4/1997	Crockett et al.	6,836,779 B2	12/2004	Poulin
5,638,509 A	6/1997	Dunphy et al.	6,847,984 B1	1/2005	Midgley et al.
5,673,381 A	9/1997	Huai et al.	6,886,020 B1	4/2005	Zahavi et al.
5,699,361 A	12/1997	Ding et al.	6,947,935 B1	9/2005	Horvitz et al.
5,729,743 A	3/1998	Squibb	6,983,322 B1	1/2006	Tripp et al.
5,737,747 A	4/1998	Vishlitzky et al.	6,996,616 B1	2/2006	Leighton et al.
5,751,997 A	5/1998	Kullick et al.	7,035,880 B1	4/2006	Crescenti et al.
5,758,359 A	5/1998	Saxon	7,047,236 B2	5/2006	Conroy et al.
5,761,677 A	6/1998	Senator et al.	7,085,787 B2	8/2006	Beier et al.
5,764,972 A	6/1998	Crouse et al.	7,103,740 B1	9/2006	Colgrove et al.
5,778,395 A	7/1998	Whiting et al.	7,130,860 B2	10/2006	Pachet et al.
5,812,398 A	9/1998	Nielsen	7,130,970 B2	10/2006	Devassy et al.
5,813,009 A	9/1998	Johnson et al.	7,181,444 B2 *	2/2007	Porter et al. 707/3
5,813,017 A	9/1998	Morris	7,197,502 B2	3/2007	Feinsmith
5,829,046 A	10/1998	Tzelnic et al.	7,240,100 B1	7/2007	Wein et al.
5,832,510 A	11/1998	Ito et al.	7,246,207 B2	7/2007	Kottomtharayil et al.
5,875,478 A	2/1999	Blumenau	7,246,211 B1	7/2007	Beloussov et al.
5,887,134 A	3/1999	Ebrahim	7,330,997 B1	2/2008	Odom
5,892,917 A	4/1999	Myerson	7,346,623 B2	3/2008	Prahlad et al.
5,901,327 A	5/1999	Ofek	7,346,676 B1	3/2008	Swildens et al.
5,907,621 A	5/1999	Bachman et al.	7,356,660 B2	4/2008	Matsunami et al.
5,924,102 A	7/1999	Perks	7,386,663 B2	6/2008	Cousins
5,950,205 A	9/1999	Aviani, Jr.	7,440,966 B2	10/2008	Adkins et al.
5,953,721 A	9/1999	Doi et al.	7,454,569 B2	11/2008	Kavuri et al.
5,974,563 A	10/1999	Beeler, Jr.	7,533,103 B2	5/2009	Brendle et al.
6,021,415 A	2/2000	Cannon et al.	7,583,861 B2	9/2009	Hanna et al.
6,026,414 A	2/2000	Anglin	2002/0069324 A1	6/2002	Gerasimov et al.
6,052,735 A	4/2000	Ulrich et al.	2002/0087550 A1	7/2002	Carlyle et al.
6,061,692 A	5/2000	Thomas et al.	2003/0115219 A1	6/2003	Chadwick
6,076,148 A	6/2000	Kedem et al.	2003/0130993 A1	7/2003	Mendelevitch et al.
6,094,416 A	7/2000	Ying	2003/0182583 A1	9/2003	Turco
6,131,095 A	10/2000	Low et al.	2004/0010493 A1	1/2004	Kojima et al.
6,131,190 A	10/2000	Sidwell	2004/0015514 A1	1/2004	Melton et al.
6,148,412 A	11/2000	Cannon et al.	2004/0254919 A1	12/2004	Giuseppini
6,154,787 A	11/2000	Urevig et al.	2004/0260678 A1	12/2004	Verbowski et al.
6,154,852 A	11/2000	Amundson et al.	2005/0050075 A1	3/2005	Okamoto et al.
6,161,111 A	12/2000	Mutalik et al.	2005/0114406 A1	5/2005	Borthakur et al.
6,167,402 A	12/2000	Yeager	2005/0154695 A1	7/2005	Gonzalez et al.
6,212,512 B1	4/2001	Barney et al.	2005/0188248 A1	8/2005	O'Brien et al.
6,260,069 B1	7/2001	Anglin	2005/0193128 A1	9/2005	Dawson et al.
6,269,431 B1	7/2001	Dunham	2005/0216453 A1	9/2005	Sasaki et al.
6,275,953 B1	8/2001	Vahalia et al.	2005/0228794 A1	10/2005	Navas et al.
6,301,592 B1	10/2001	Aoyama et al.	2005/0262097 A1	11/2005	Sim-Tang et al.
6,324,581 B1	11/2001	Xu et al.	2006/0004820 A1	1/2006	Claudatos et al.
6,328,766 B1	12/2001	Long	2006/0010227 A1	1/2006	Atluri
6,330,570 B1	12/2001	Crighton et al.	2006/0031225 A1	2/2006	Palmeri et al.
6,330,642 B1	12/2001	Carteau	2006/0101285 A1	5/2006	Chen et al.
6,343,324 B1	1/2002	Hubis et al.	2006/0106814 A1	5/2006	Blumenau et al.
RE37,601 E	3/2002	Eastridge et al.	2006/0253495 A1 *	11/2006	Png 707/200
6,356,801 B1	3/2002	Goodman et al.	2006/0259468 A1	11/2006	Brooks et al.
6,374,336 B1	4/2002	Peters et al.	2006/0294094 A1	12/2006	King et al.
6,389,432 B1	5/2002	Pothapragada et al.	2007/0033191 A1 *	2/2007	Hornkvist et al. 707/9
6,418,478 B1	7/2002	Ignatius et al.	2007/0112809 A1 *	5/2007	Arrouye et al. 707/100
6,421,683 B1	7/2002	Lamburt			

2008/0091655 A1 4/2008 Gokhale et al.

FOREIGN PATENT DOCUMENTS

EP	0405926	1/1991
EP	0467546	1/1992
EP	0774715	5/1997
EP	0809184	11/1997
EP	0899662	3/1999
EP	0981090	2/2000
EP	1174795	1/2002
WO	WO-95/13580	5/1995
WO	WO-99/12098	3/1999
WO	WO-99/14692	3/1999
WO	WO-03060774	7/2003
WO	WO-2005/055093	6/2005
WO	PCT/US2006/045556	11/2006
WO	WO-2007062254	5/2007
WO	WO-2007062429	5/2007
WO	WO-2008049023	4/2008

OTHER PUBLICATIONS

PCT International Search Report and Written Opinion for International Application No. PCT/US07/81681, Mail Date Oct. 20, 2008, 11 pages.

U.S. Appl. No. 11/931,034, filed Oct. 31, 2007, Kottomtharayil et al.

U.S. Appl. No. 12/058,487, filed Mar. 28, 2008, Prahlad.

U.S. Appl. No. 12/058,575, filed Mar. 28, 2008, Prahlad et al.

U.S. Appl. No. 12/058,589, filed Mar. 28, 2008, Prahlad et al.

International Search Report and Written Opinion, PCT Application No. PCT/US2006/61304, Applicant: Commvault Systems, Inc, Date of Mailing: Mar. 17, 2008, 17 pages.

EMC Corporation, "Today's Choices for Business Continuity," 2004, 12 pages.

Microsoft Developer Network, "Win32_File_Attribute_Data," online library article, [accessed on Nov. 10, 2005], 3 pages.

Microsoft Developer Network, "GetFileAttributesEx," online library article, [accessed on Nov. 10, 2005], 2 pages.

Microsoft Developer Network, "GetFileAttributes," online library article, [accessed on Nov. 10, 2005], 3 pages.

Jeffrey Richter and Luis Felipe Cabrera, "A File System for the 21st Century: Previewing the Windows NT 5.0 File System," and attached text figures, Microsoft Systems Journal, Nov. 1998, 24 pages.

Jeffrey Cooperstein and Jeffrey Richter, "Keeping an Eye on Your NTFS Drives: the Windows 2000 Change Journal Explained," Microsoft Systems Journal, Sep. 1999, 17 pages.

Jeffrey Cooperstein and Jeffrey Richter, "Keeping an Eye on Your NTFS Drives, Part II: Building a Change Journal Application," Microsoft Systems Journal, Oct. 1999, 14 pages.

Brad O'Neill, "New Tools to Classify Data," Storage Magazine, Aug. 2005, 4 pages.

Karl Langdon and John Merryman, "Data Classification: Getting Started," Storage Magazine, Jul. 2005, 3 pages.

U.S. Appl. No. 11/563,940, Prahlad et al.

U.S. Appl. No. 11/564,119, Prahlad et al.

U.S. Appl. No. 11/564,130, Prahlad et al.

U.S. Appl. No. 11/564,136, Prahlad et al.

U.S. Appl. No. 11/564,163, Prahlad et al.

U.S. Appl. No. 11/564,170, Prahlad et al.

U.S. Appl. No. 11/564,174, Prahlad et al.

U.S. Appl. No. 11/564,177, Prahlad et al.

U.S. Appl. No. 11/564,180, Prahlad et al.

U.S. Appl. No. 11/564,194, Prahlad et al.

U.S. Appl. No. 11/564,197, Prahlad et al.

U.S. Appl. No. 11/564,215, Prahlad et al.

U.S. Appl. No. 11/564,220, Prahlad et al.

U.S. Appl. No. 11/564,221, Prahlad et al.

U.S. Appl. No. 11/564,233, Prahlad et al.

U.S. Appl. No. 11/605,931, Prahlad et al.

U.S. Appl. No. 11/605,932, Prahlad et al.

U.S. Appl. No. 11/605,944, Prahlad et al.

Armstead et al., "Implementation of a Campus-wide Distributed Mass Storage Service: The Dream vs. Reality," *IEEE*, 1995, pp. 190-199.

Arneson, "Mass Storage Archiving in Network Environments," Digest of Papers, Ninth IEEE Symposium on Mass Storage Systems, Oct. 31, 1988-Nov. 3, 1988, pp. 45-50, Monterey, CA.

Cabrera et al., "ADSM: A Multi-Platform, Scalable, Backup and Archive Mass Storage System," Digest of Papers, Compton '95, Proceedings of the 40th IEEE Computer Society International Conference, Mar. 5, 1995-Mar. 9, 1995, pp. 420-427, San Francisco, CA.

Eitel, "Backup and Storage Management in Distributed Heterogeneous Environments," *IEEE*, 1994, pp. 124-126.

Jander, M., "Launching Storage-Area Net," *Data Communications*, US, McGraw Hill, NY, vol. 27, No. 4 (Mar. 21, 1998), pp. 64-72.

Jason Gait, "The Optical File Cabinet: A Random-Access File System For Write-Once Optical Disks," *IEEE Computer*, vol. 21, No. 6, pp. 11-22 (1988) (see in particular figure 5 in p. 15 and recitation in claim 5).

Rosenblum et al., "The Design and Implementation of a Log-Structured File System," *Operating Systems Review SIGOPS*, vol. 25, No. 5, New York, US, pp. 1-15 (May 1991).

Supplementary European Search Report for European Application EP06846386, Dated Dec. 30, 2008, European Patent Office, 6 pages.

Final Office Action for U.S. Appl. No. 11/931,034, Mail Date Dec. 29, 2008, 18 pages.

Partial International Search Results, mailed May 25, 2007, International Application No. PCT/US2006/045556, 2 pages.

"Text Figures", retrieved from <http://www.microsoft.com/msj/1198.ntfs/ntfstextfigs.htm> on Nov. 10, 2005, 7 pages.

European Patent Office Examination Report for EP application 06846386.8, Mail Date Apr. 29, 2009, 6 pages.

Non-Final Office Action for U.S. Appl. No. 11/564,215, Mail Date May 8, 2009, 39 pages.

Non-Final Office Action for U.S. Appl. No. 11/564,136, Mail Date May 15, 2009, 25 pages.

U.S. Appl. No. 12/511,653, filed Jul. 29, 2009, Prahlad et al.

Communication with extended European Search Report for Application No. PCT/US2006/061304, dated Dec. 30, 2008.

* cited by examiner

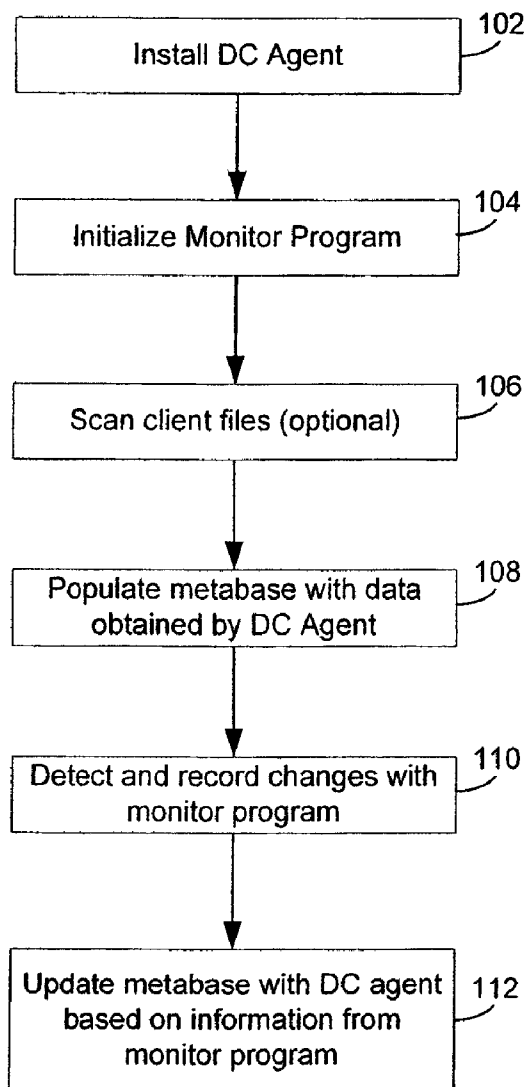
100

FIG. 1

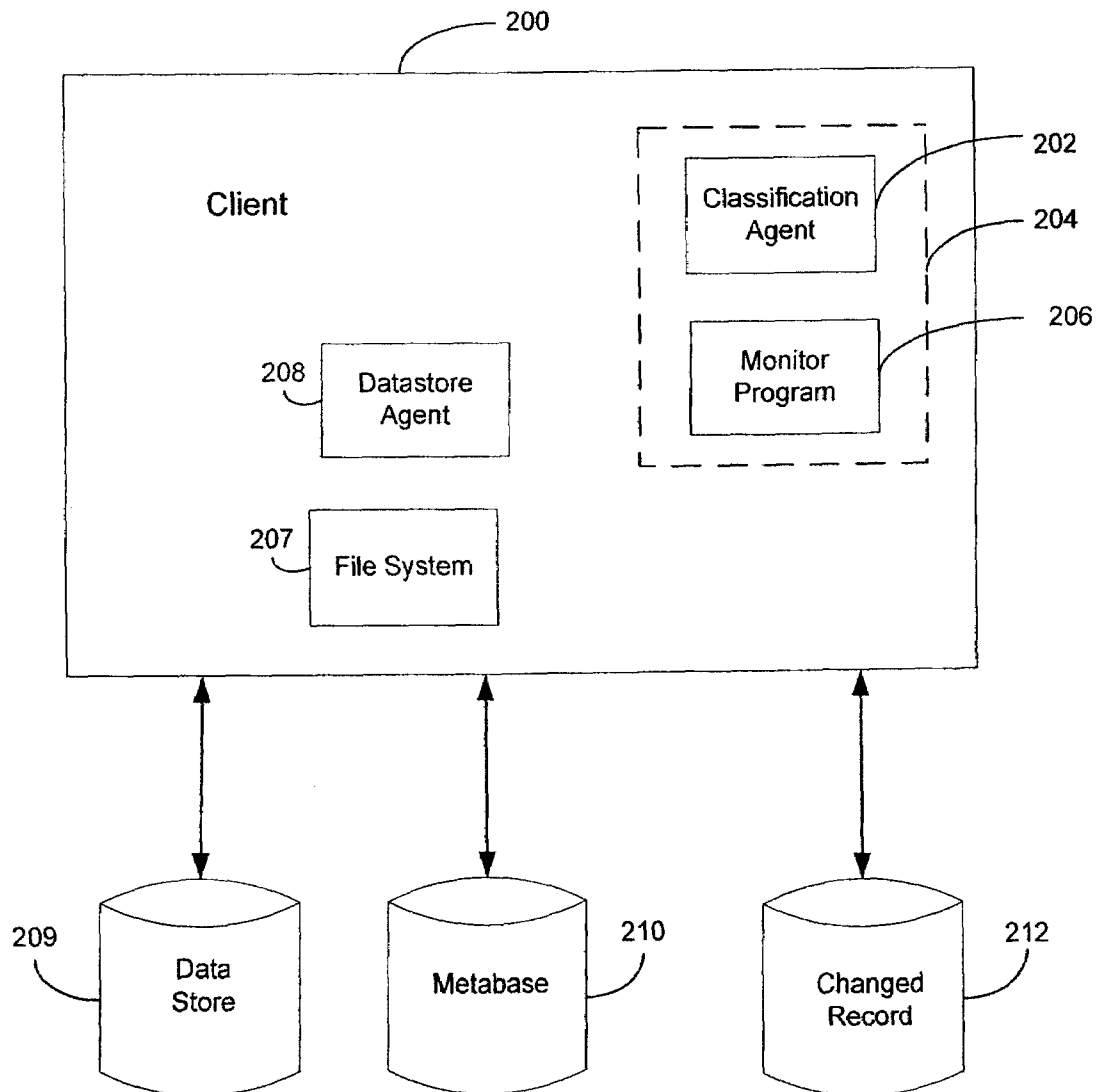


FIG. 2

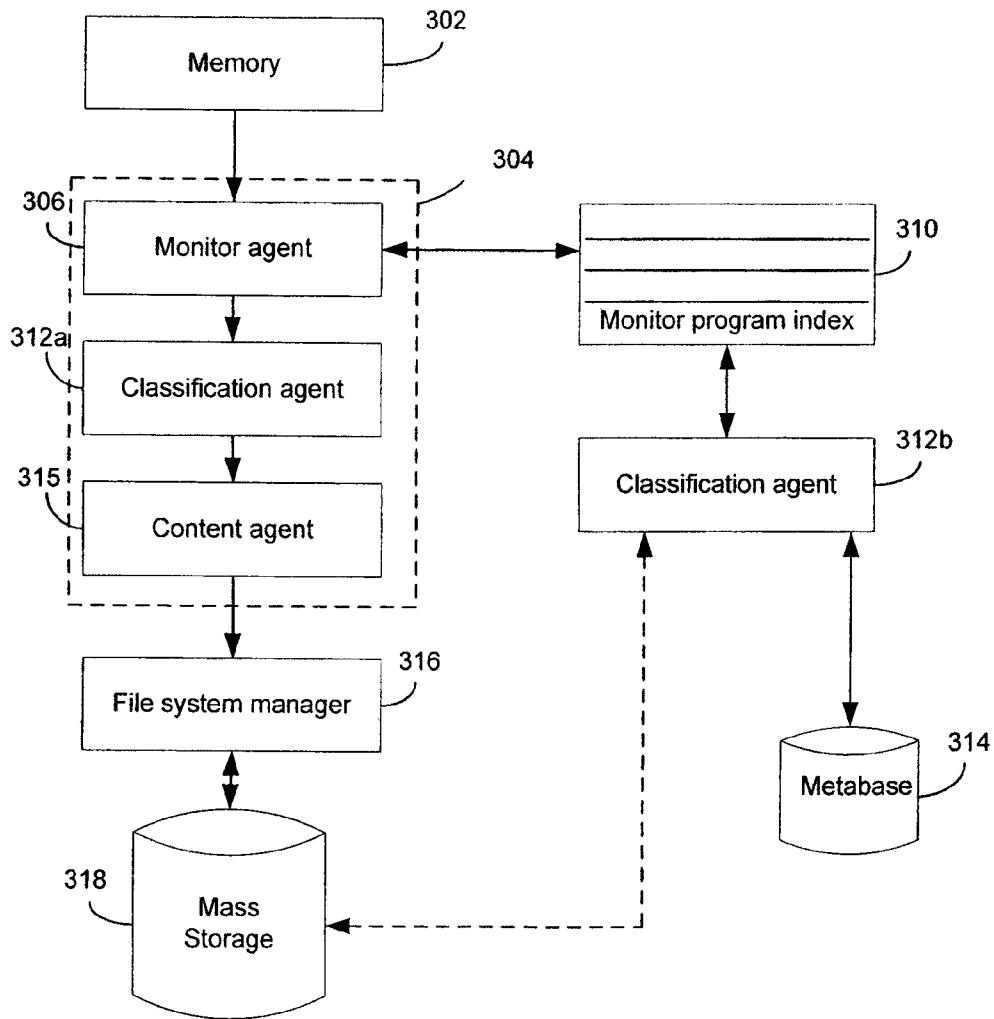
300

FIG. 3a

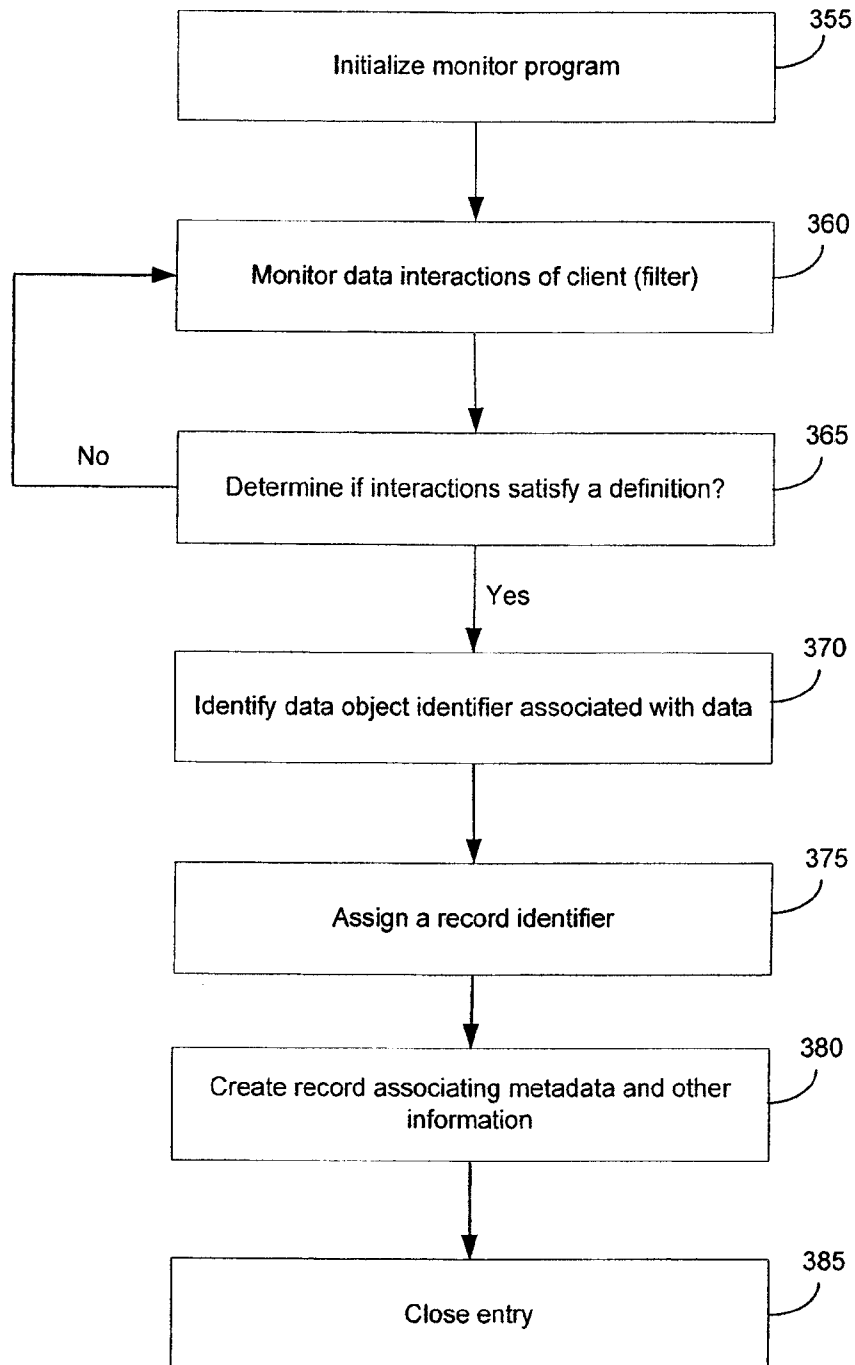
350

FIG. 3b

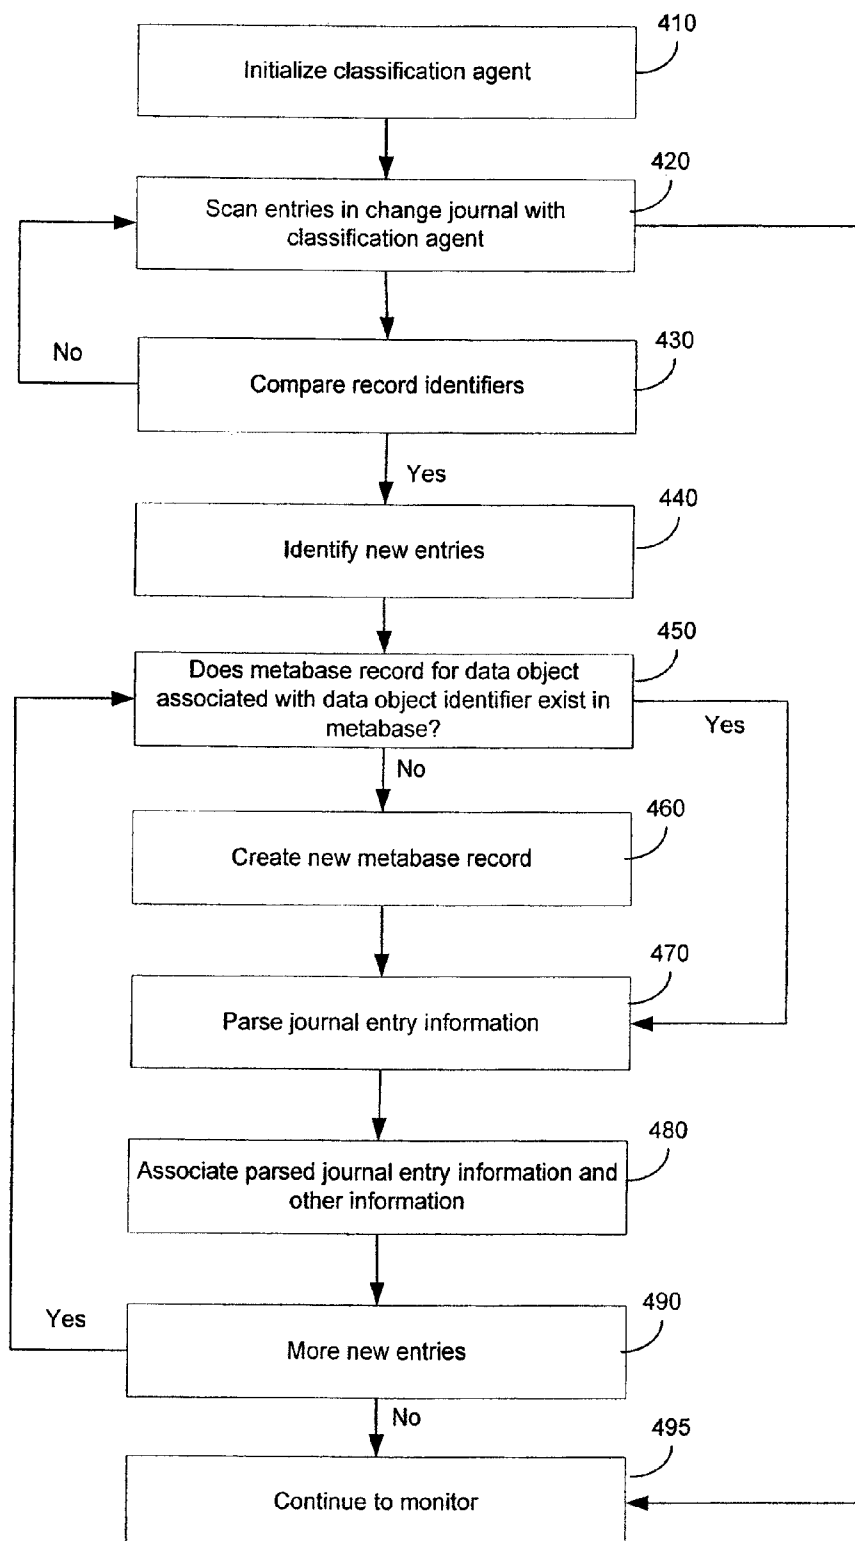
400

FIG. 4

500




FIG. 5

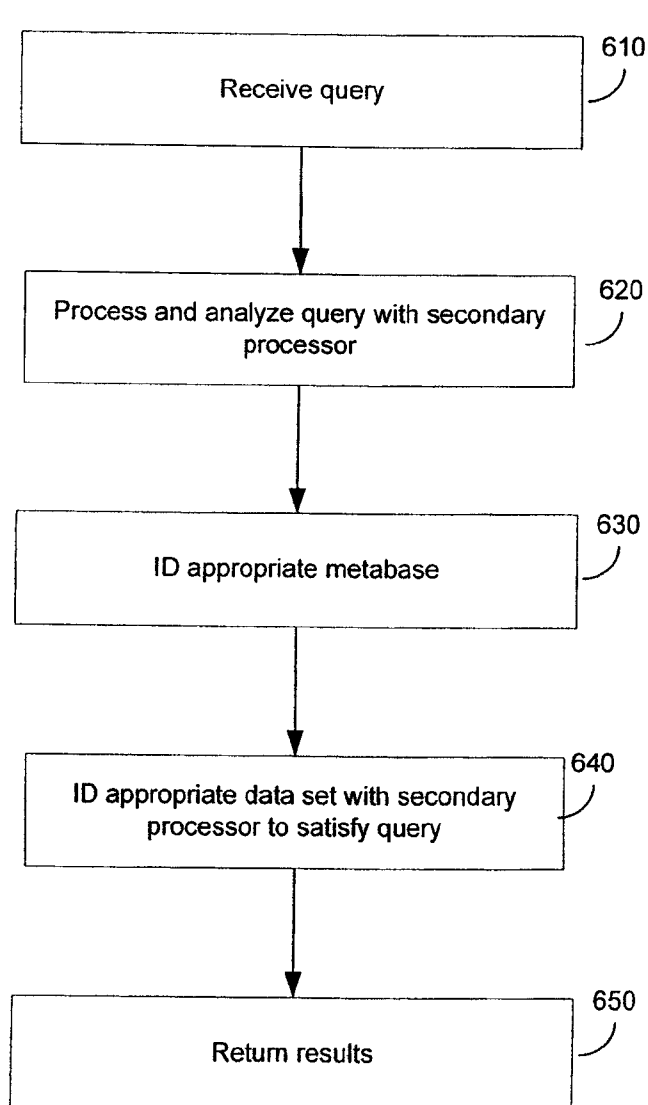


FIG. 6

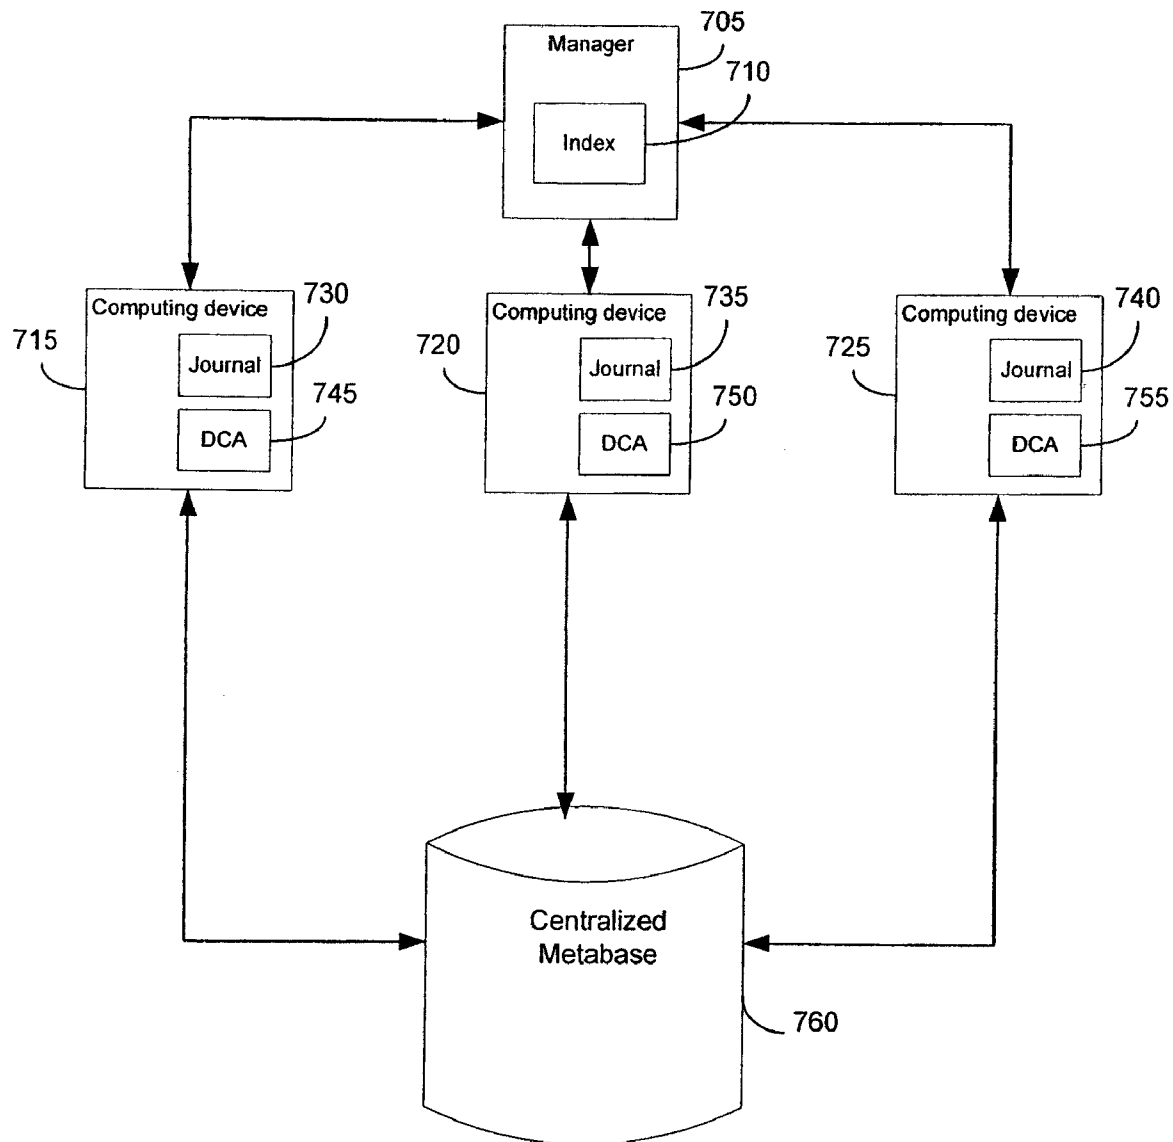
700

FIG. 7

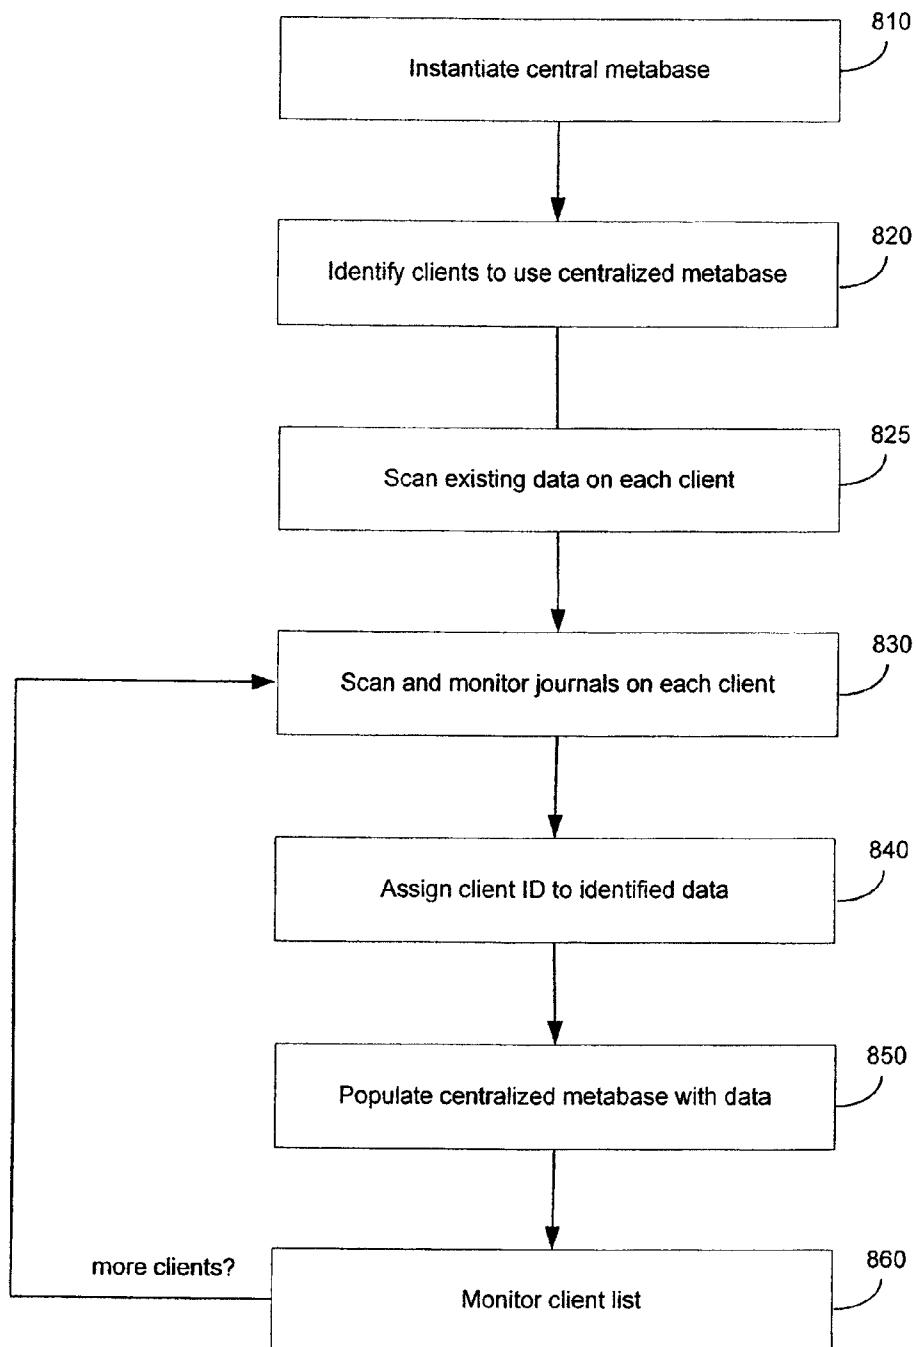
800

FIG. 8

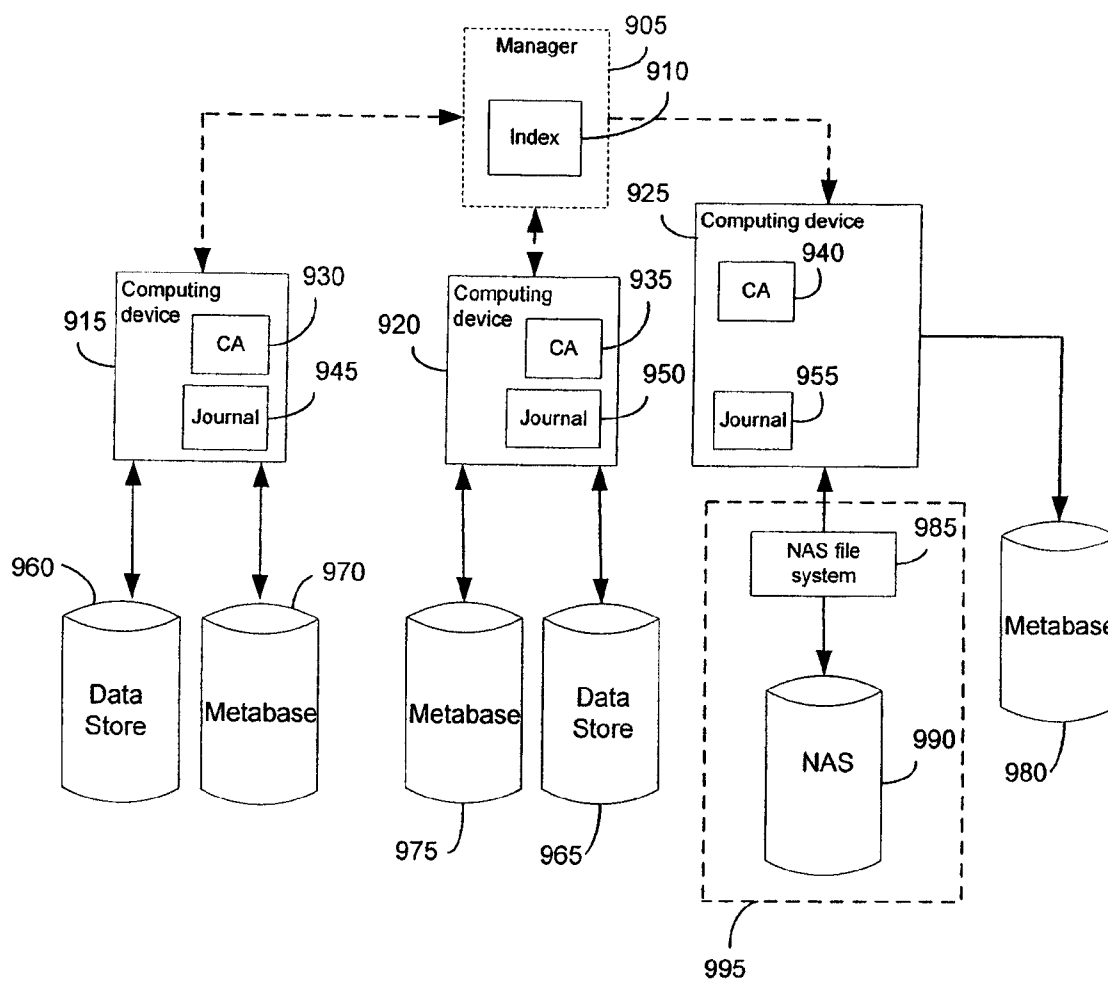
900

FIG. 9

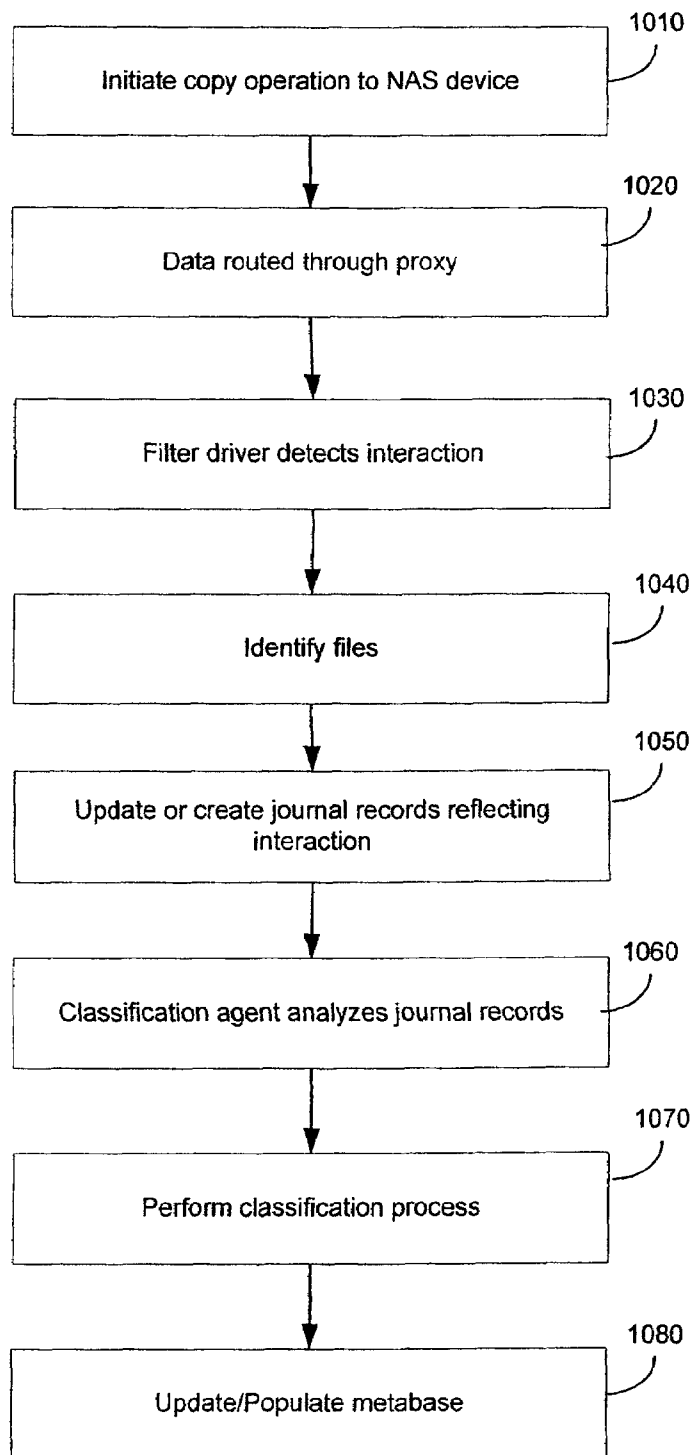
1000

FIG. 10

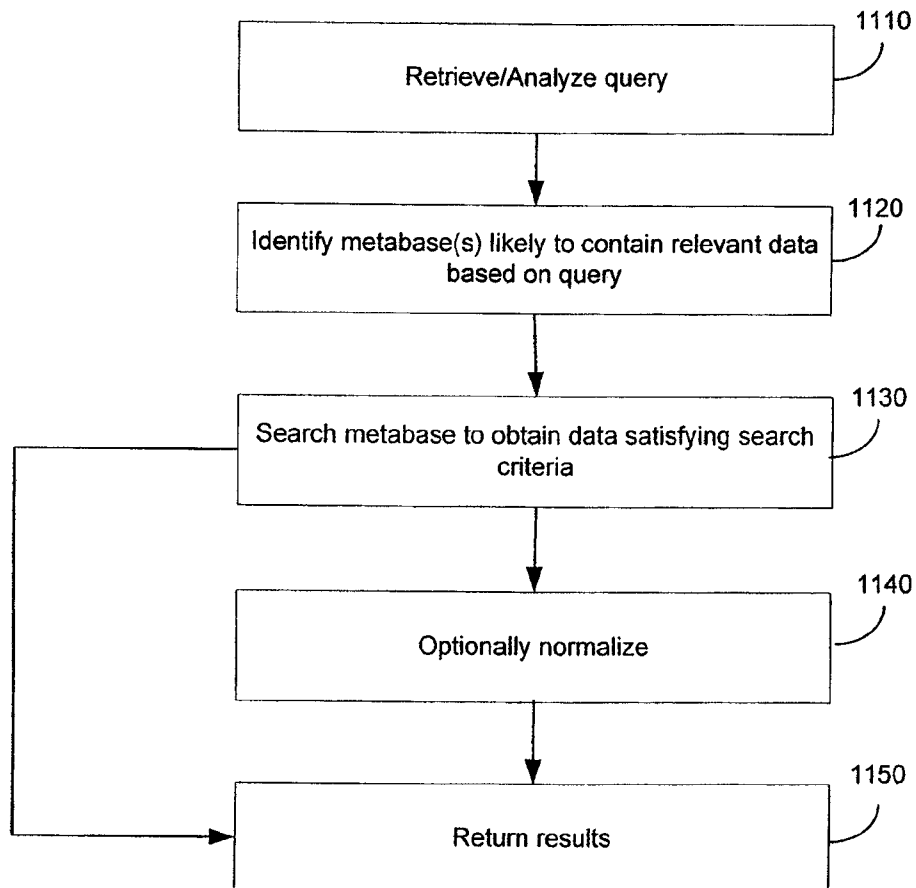
1100

FIG. 11

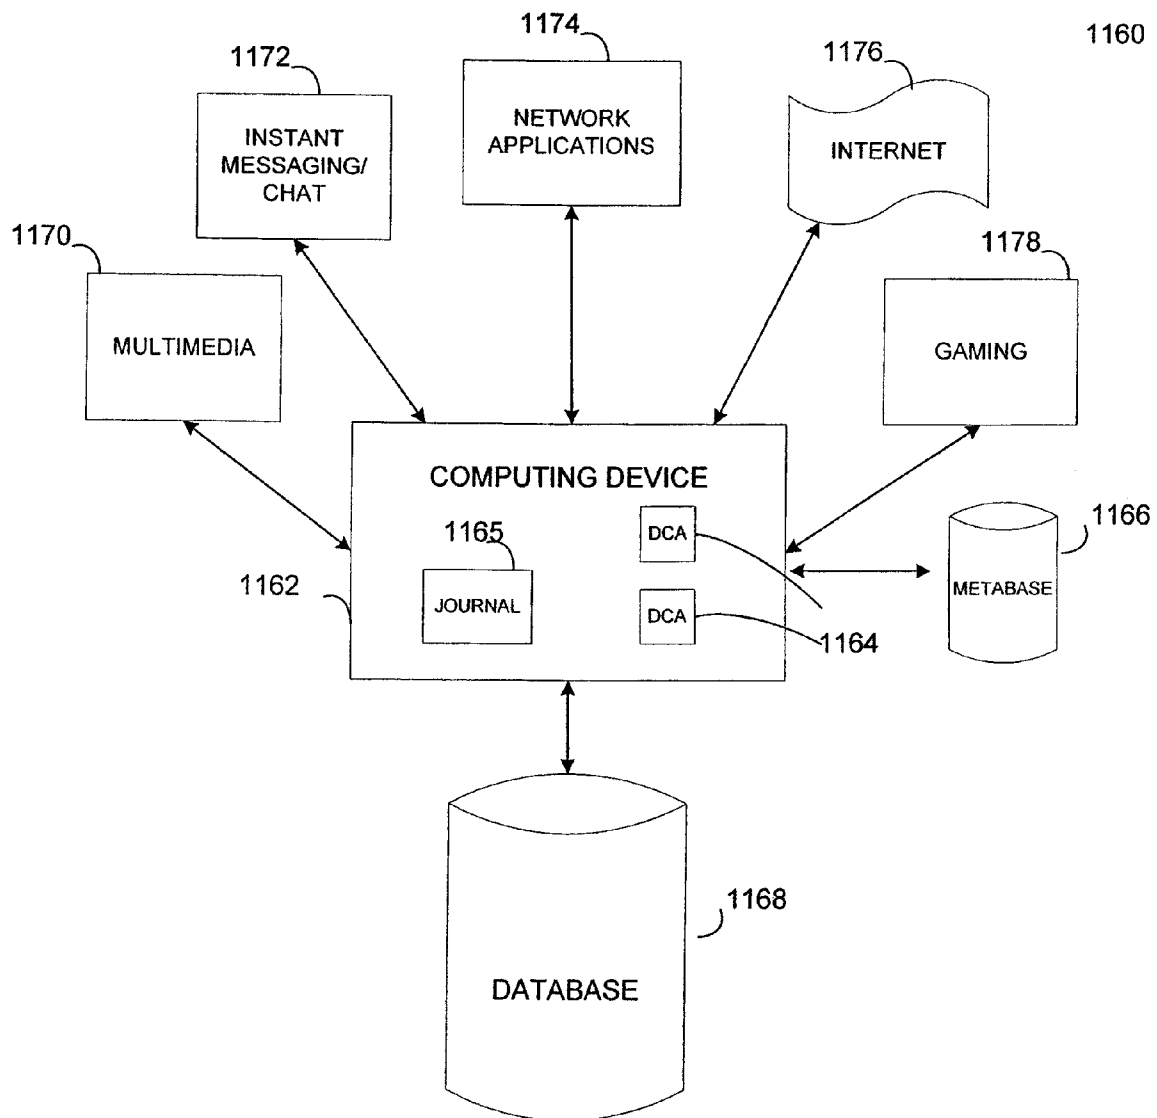


Figure 11a

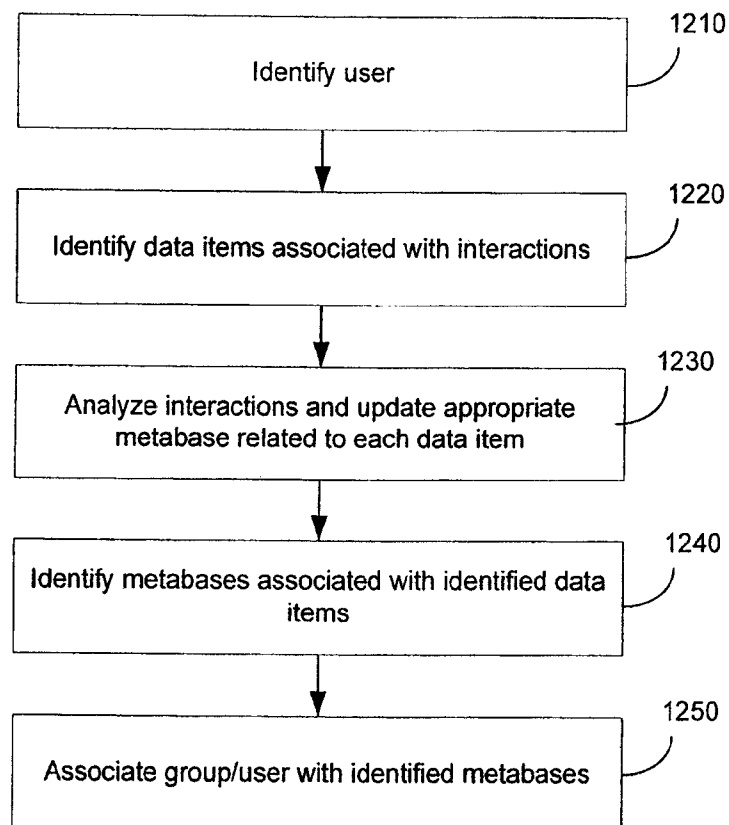
1200

FIG. 12

1300




FIG. 13

1400

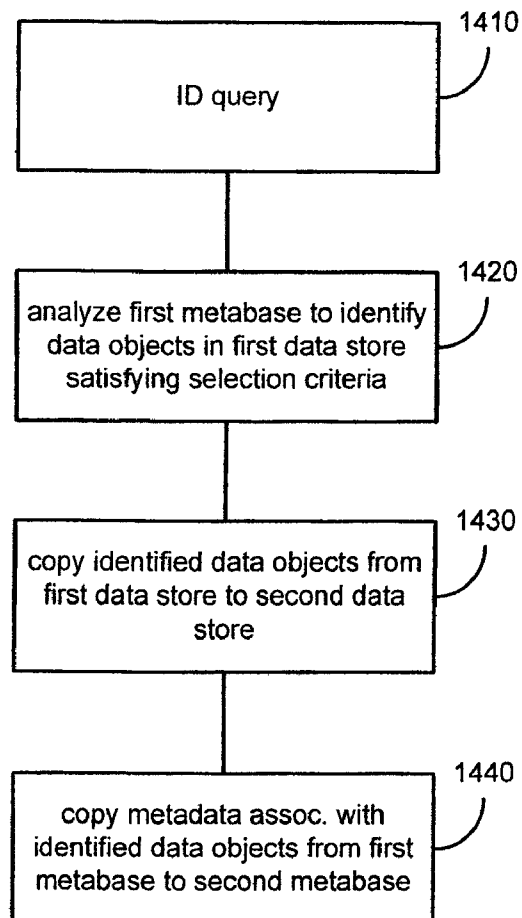


FIG. 14




FIG. 15

1600


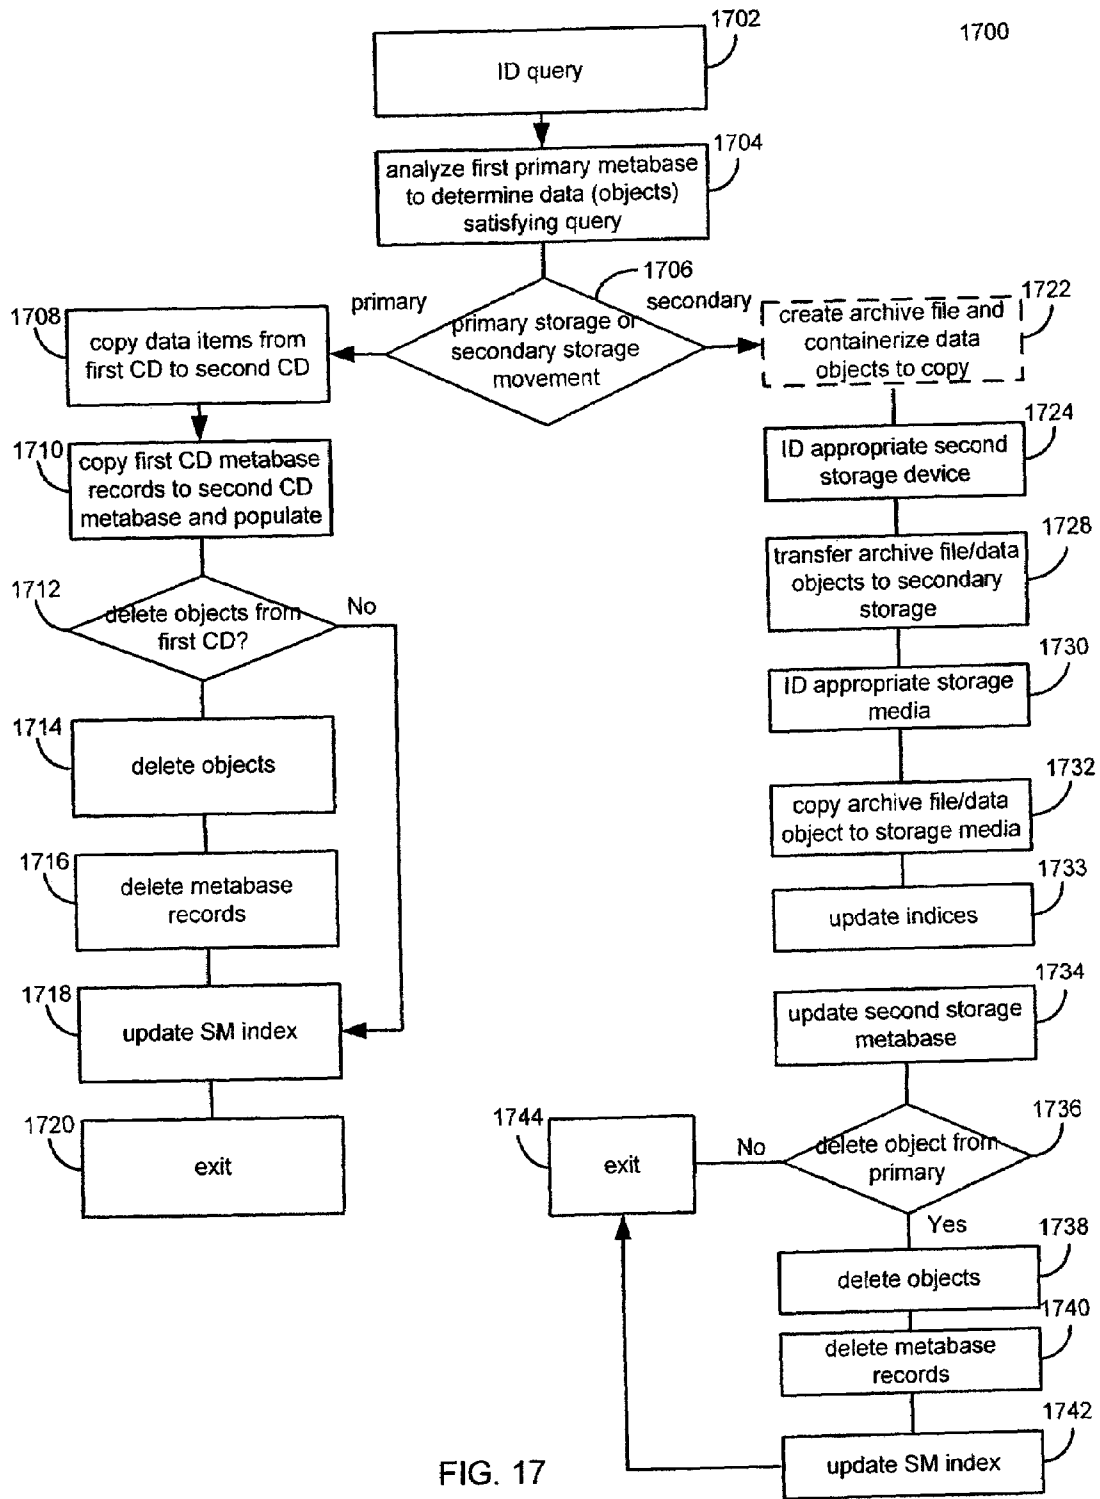


FIG. 16



1800

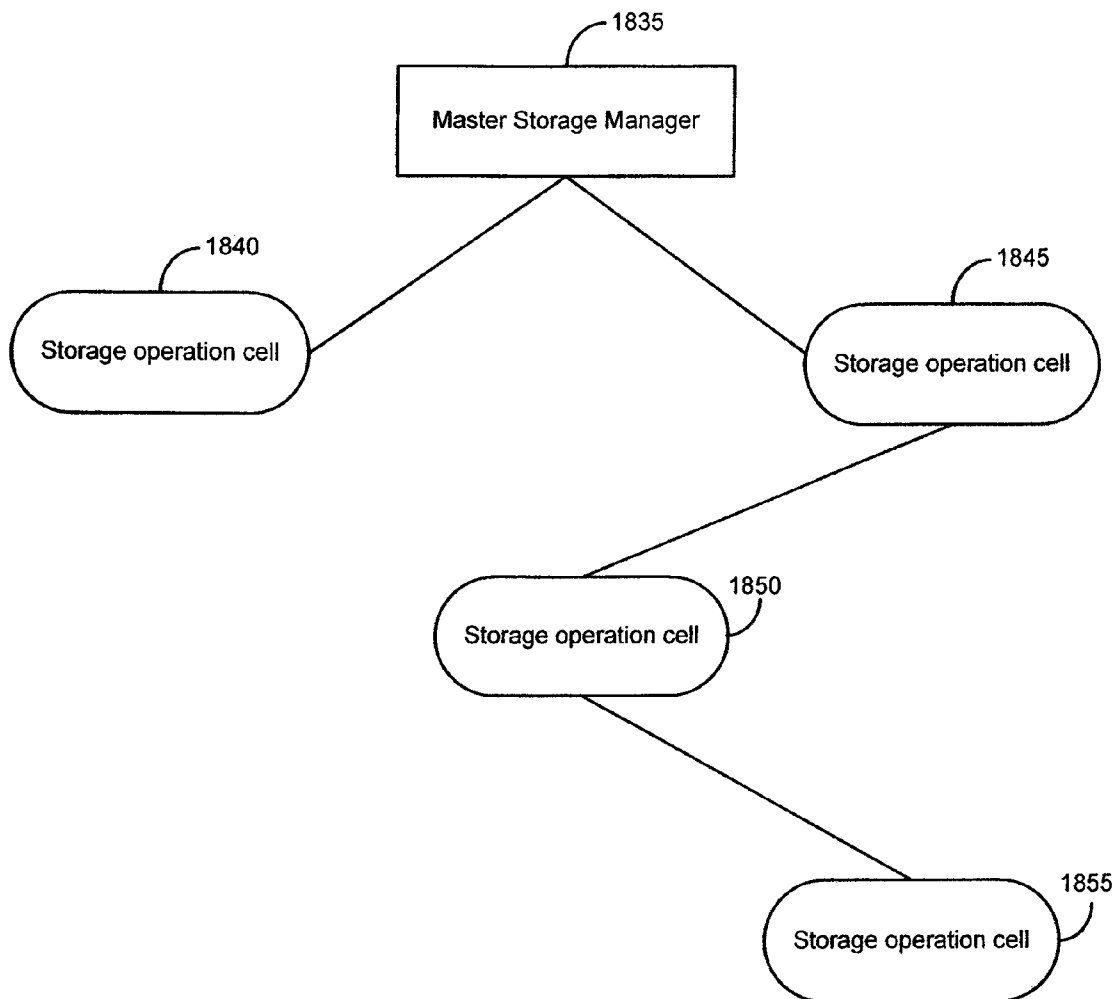


FIG. 18

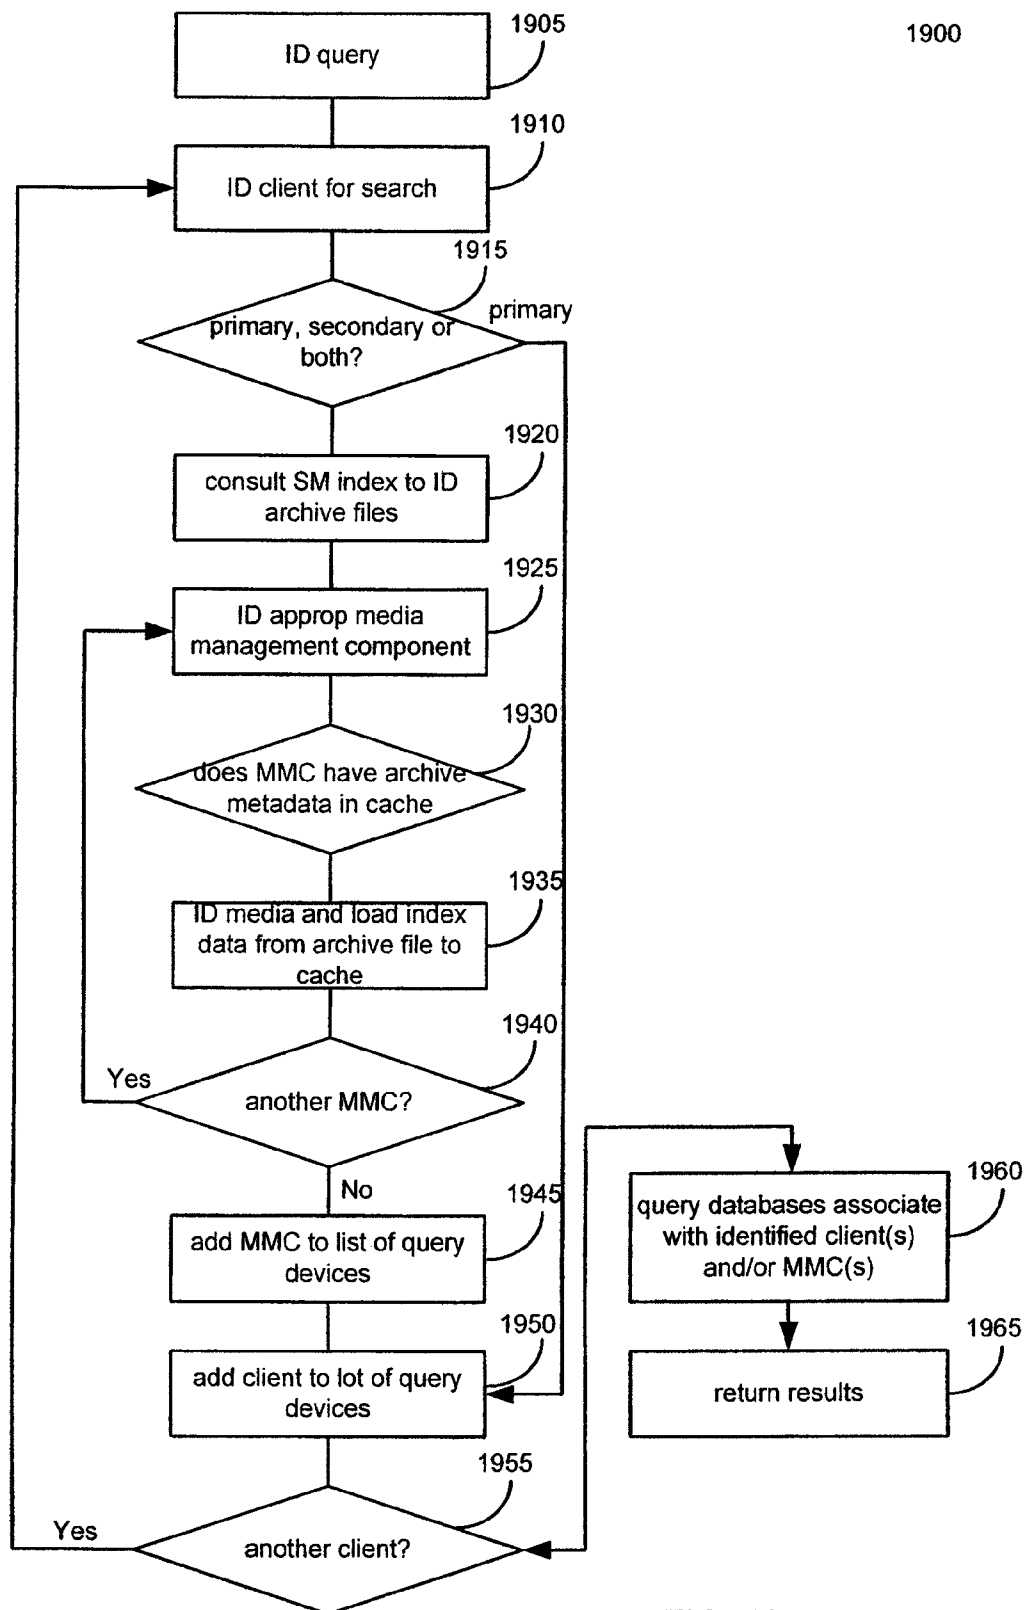


FIG. 19

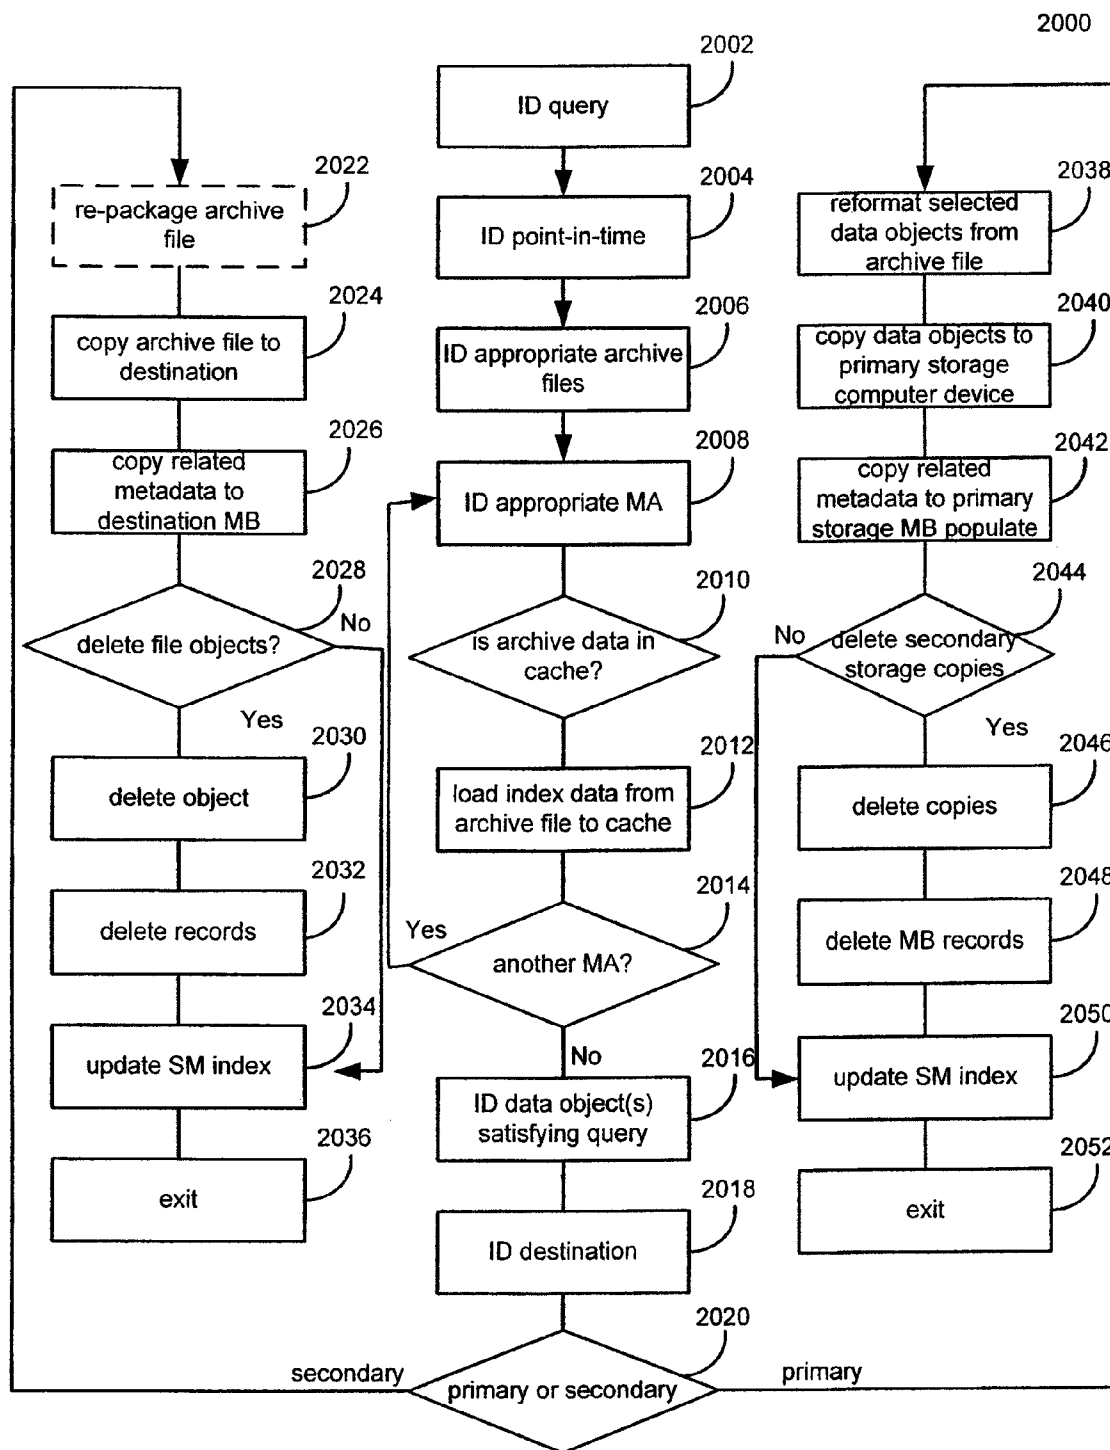


FIG. 20

1

SYSTEMS AND METHODS FOR CLASSIFYING AND TRANSFERRING INFORMATION IN A STORAGE NETWORK

CROSS-REFERENCE TO RELATED APPLICATION(S)

This application claims the benefit of U.S. Provisional Application No. 60/740,686, entitled SYSTEMS AND METHODS FOR CLASSIFYING INFORMATION IN A STORAGE NETWORK, filed Nov. 28, 2005 and U.S. Provisional Application No. 60/752,203, entitled SYSTEMS AND METHODS FOR CLASSIFYING AND TRANSFERRING INFORMATION IN A STORAGE NETWORK, filed Dec. 19, 2005, both of which are hereby incorporated herein by reference in their entirety.

This application is related to the following patents and pending U.S. applications, each of which is hereby incorporated herein by reference in its entirety:

application Ser. No. 09/354,058, titled HIERARCHICAL BACKUP AND RETRIEVAL SYSTEM, filed Jul. 15, 1999;

U.S. Pat. No. 6,418,478, titled PIPELINED HIGH SPEED DATA TRANSFER MECHANISM, issued Jul. 9, 2002; application Ser. No. 09/610,738, titled MODULAR BACKUP AND RETRIEVAL SYSTEM USED IN CONJUNCTION WITH A STORAGE AREA NETWORK, filed Jul. 6, 2000;

U.S. Pat. No. 6,542,972 titled LOGICAL VIEW AND ACCESS TO PHYSICAL STORAGE IN MODULAR DATA AND STORAGE MANAGEMENT SYSTEM, issued Apr. 1, 2003;

U.S. Pat. No. 6,658,436, titled LOGICAL VIEW AND ACCESS TO DATA MANAGE BY A MODULAR DATA AND STORAGE MANAGEMENT SYSTEM, issued Dec. 2, 2003;

application Ser. No. 10/658,095, titled DYNAMIC STORAGE DEVICE POOLING IN A COMPUTER SYSTEM, filed Sep. 9, 2003;

application Ser. No. 10/262,556, titled METHOD FOR MANAGING SNAPSHOTS GENERATED BY AN OPERATING SYSTEM OR OTHER APPLICATION, filed Sep. 30, 2002;

application Ser. No. 10/818,749, SYSTEM AND METHOD FOR DYNAMICALLY PERFORMING STORAGE OPERATIONS IN A COMPUTER NETWORK, filed Apr. 5, 2004;

application Ser. No. 10/877,831, HIERARCHICAL SYSTEM AND METHOD FOR PERFORMING STORAGE OPERATIONS IN A COMPUTER NETWORK, filed Jun. 25, 2004;

application Ser. No. 11/313,256 entitled SYSTEM AND METHOD FOR CONTAINERIZED DATA STORAGE AND TRACKING, filed Dec. 19, 2005,

Application Ser. No. 60/752,198, entitled "Systems and Methods for Granular Resource Management in a Storage Network" filed Dec. 19, 2005,

application Ser. No. 11/313,224, entitled "Systems and Methods for Performing Multi-Path Storage Operations" filed Dec. 19, 2005,

Application Ser. No. 60/752,196, entitled "Systems and Methods for Migrating Components in a Hierarchical Storage Network" filed Dec. 19, 2005,

Application Ser. No. 60/752,202, entitled "Systems and Methods for Unified Reconstruction of Data in a Storage Network" filed Dec. 19, 2005,

2

Application Ser. No. 60/752,201, entitled "Systems and Methods for Resynchronizing Storage Operations" filed Dec. 19, 2005,

Application Ser. No. 60/752,197, entitled "Systems and Methods for Hierarchical Client Group Management" filed Dec. 19, 2005,

BACKGROUND

Aspects of the invention disclosed herein relate generally to performing operations on electronic data in a computer network. More particularly, aspects of the present invention relate to detecting data interactions within a computer network and/or performing storage-related operations for a computer network according to a specified classification paradigm.

Current storage management systems employ a number of different methods to perform storage operations on electronic data. For example, data can be stored in primary storage as a primary copy or in secondary storage as various types of secondary copies including, as a backup copy, a snapshot copy, a hierarchical storage management copy ("HSM"), as an archive copy, and as other types of copies.

A primary copy of data is generally a production copy or other "live" version of the data which is used by a software application and is generally in the native format of that application. Primary copy data may be maintained in a local memory or other high-speed storage device that allows for relatively fast data access if necessary. Such primary copy data is typically intended for short term retention (e.g., several hours or days) before some or all of the data is stored as one or more secondary copies, for example to prevent loss of data in the event a problem occurred with the data stored in primary storage.

Secondary copies include point-in-time data and are typically for intended for long-term retention (e.g., weeks, months or years depending on retention criteria, for example as specified in a storage policy as further described herein) before some or all of the data is moved to other storage or discarded. Secondary copies may be indexed so users can browse and restore the data at another point in time. After certain primary copy data is backed up, a pointer or other location indicia such as a stub may be placed in the primary copy to indicate the current location of that data.

One type of secondary copy is a backup copy. A backup copy is generally a point-in-time copy of the primary copy data stored in a backup format as opposed to in native application format. For example, a backup copy may be stored in a backup format that is optimized for compression and efficient long-term storage. Backup copies generally have relatively long retention periods and may be stored on media with slower retrieval times than other types of secondary copies and media. In some cases, backup copies may be stored at an offsite location.

Another form of secondary copy is a snapshot copy. From an end-user viewpoint, a snapshot may be thought of as an instant image of the primary copy data at a given point in time. A snapshot generally captures the directory structure of a primary copy volume at a particular moment in time, and also preserves file attributes and contents. In some embodiments, a snapshot may exist as a virtual file system, parallel to the actual file system. Users typically gain a read-only access to the record of files and directories of the snapshot. By electing to restore primary copy data from a snapshot taken at a given point in time, users may also return the current file system to the prior state of the file system that existed when the snapshot was taken.

A snapshot may be created instantly, using a minimum of file space, but may still function as a conventional file system backup. A snapshot may not actually create another physical copy of all the data, but may simply create pointers that are able to map files and directories to specific disk blocks.

In some embodiments, once a snapshot has been taken, subsequent changes to the file system typically do not overwrite the blocks in use at the time of snapshot. Therefore, the initial snapshot may use only a small amount of disk space needed to record a mapping or other data structure representing or otherwise tracking the blocks that correspond to the current state of the file system. Additional disk space is usually only required when files and directories are actually modified later. Furthermore, when files are modified, typically only the pointers which map to blocks are copied, not the blocks themselves. In some embodiments, for example in the case of copy-on-write snapshots, when a block changes in primary storage, the block is copied to secondary storage before the block is overwritten in primary storage and the snapshot mapping of file system data is updated to reflect the changed block(s) at that particular point in time. An HSM copy is generally a copy of the primary copy data, but typically includes only a subset of the primary copy data that meets a certain criteria and is usually stored in a format other than the native application format. For example, an HSM copy might include only that data from the primary copy that is larger than a given size threshold or older than a given age threshold and that is stored in a backup format. Often, HSM data is removed from the primary copy, and a stub is stored in the primary copy to indicate its new location. When a user requests access to the HSM data that has been removed or migrated, systems use the stub to locate the data and often make recovery of the data appear transparent even though the HSM data may be stored at a location different from the remaining primary copy data.

An archive copy is generally similar to an HSM copy, however, the data satisfying criteria for removal from the primary copy is generally completely removed with no stub left in the primary copy to indicate the new location (i.e., where it has been moved to). Archive copies of data are generally stored in a backup format or other non-native application format. In addition, archive copies are generally retained for very long periods of time (e.g., years) and in some cases are never deleted. Such archive copies may be made and kept for extended periods in order to meet compliance regulations or for other permanent storage applications.

In some embodiments, application data over its lifetime moves from more expensive quick access storage to less expensive slower access storage. This process of moving data through these various tiers of storage is sometimes referred to as information lifecycle management ("ILM"). This is the process by which data is "aged" from more forms of secondary storage with faster access/restore times down through less expensive secondary storage with slower access/restore times, for example, as the data becomes less important or mission critical over time.

Examples of various types of data and copies of data are further described in the above-referenced related applications that are hereby incorporated by reference in their entirety. One example of a system that performs storage operations on electronic data that produce such copies is the QiNetix storage management system by CommVault Systems of Oceanport, N.J.

The QiNetix system leverages a modular storage management architecture that may include, among other things, storage manager components, client or data agent components, and media agent components as further described in U.S.

patent application Ser. No. 10/818,749 which is hereby incorporated herein by reference in its entirety. The QiNetix system also may be hierarchically configured into backup cells to store and retrieve backup copies of electronic data as further described in U.S. patent application Ser. No. 09/354,058 which is hereby incorporated by reference in its entirety.

Regardless of where data is stored, conventional storage management systems perform storage operations associated with electronic data based on location-specific criteria. For example, data generated by applications running on a particular client is typically copied according to location-specific criteria, such as from a certain location such as a specific folder or subfolder, according to a specified data path, etc. A module installed on the client or elsewhere in the system may supervise the transfer of data from the client to another location in a primary or secondary storage. Similar data transfers associated with location-specific criteria are performed when restoring data from secondary storage to primary storage. For example, to restore data, a user or system process must specify a particular secondary storage device, piece of media, archive file, etc. Thus, the precision with which conventional storage management systems perform storage operations on electronic data is generally limited by the ability to define or specify storage operations based on data location rather than information relating to or describing the data itself.

Moreover, when identifying data objects, such as files associated with performing storage operations, conventional storage systems often scan the file system of a client or other computing device to determine which data objects on the client should be associated with the storage operation. This may involve collecting file and/or folder attributes by traversing the file system of the client prior to performing storage operations. This process is typically time-consuming and uses significant client resources that might be more desirably spent performing other tasks associated with production applications. There is thus a need for systems and methods for performing more precise and efficient storage operations.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention is illustrated in the figures of the accompanying drawings which are meant to be exemplary and not limiting, in which like references are intended to refer to like or corresponding parts, and in which:

FIG. 1 is a flow chart in accordance with an embodiment of the present invention;

FIG. 2 is a system constructed in accordance with an embodiment of the present invention;

FIG. 3a is a system constructed in accordance with an embodiment of the present invention;

FIG. 3b is a flow chart in accordance with an embodiment of the present invention;

FIG. 4 is a flow chart in accordance with an embodiment of the present invention;

FIG. 5 is a system constructed in accordance with an embodiment of the present invention;

FIG. 6 is a flow chart in accordance with an embodiment of the present invention;

FIG. 7 is a system constructed in accordance with an embodiment of the present invention;

FIG. 8 is a flow chart in accordance with an embodiment of the present invention;

FIG. 9 is a system constructed in accordance with an embodiment of the present invention;

FIG. 10 is a flow chart in accordance with an embodiment of the present invention;

5

FIG. 11 is a flow chart in accordance with an embodiment of the present invention;

FIG. 11a is a system constructed in accordance with an embodiment of the present invention;

FIG. 12 is a flow chart in accordance with an embodiment of the present invention;

FIG. 13 is a system constructed in accordance with an embodiment of the present invention;

FIG. 14 is a flow chart in accordance with an embodiment of the present invention;

FIG. 15 is a system constructed in accordance with an embodiment of the present invention;

FIG. 16 is a data arrangement in accordance with an embodiment of the present invention;

FIG. 17 is a flow chart in accordance with an embodiment of the present invention;

FIG. 18 is a system constructed in accordance with an embodiment of the present invention;

FIG. 19 is a flow chart in accordance with an embodiment of the present invention; and

FIG. 20 is a flow chart in accordance with an embodiment of the present invention.

COPYRIGHT NOTICE

A portion of the disclosure of this patent document contains material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosures, as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all copyright rights whatsoever.

DETAILED DESCRIPTION

Aspects of the present invention are generally concerned with systems and methods that analyze and classify various forms of data that, among other things, facilitates identification, searching, storage and retrieval of data that satisfies certain criteria. Although described in connection with certain specific embodiments, it will be understood that the inventions disclosed herein have broad-based applicability to any wireless or hard-wired network or data transfer system that stores and conveys data from one point to another, including communication networks, enterprise networks, storage networks, and the like.

Aspects of the present invention provide systems and methods for data classification to facilitate and improve data management within an enterprise. The disclosed systems and methods evaluate and define data management operations based on data characteristics rather than data location, among other things. Also provided are methods for generating a data structure of metadata that describes system data and storage operations. This data structure may be consulted to determine changes in system data rather than scanning the data files themselves.

Generally, the systems and methods described in detail below are for analyzing data and other information in a computer network (sometimes referred to herein as a "data object") and creating a database or index of information which may describe certain pertinent aspects of the data objects that allow a user or system process to consult the database to obtain information regarding the network data. For example, a data collection agent may traverse a network file system and obtain certain characteristics and other attributes of data in the system. In some embodiments, such a database may be a collection of metadata and/or other infor-

6

mation regarding the network data and may be referred to herein as a "metabase." Generally, metadata refers to data or information about data, and may include, for example, data relating to storage operations or storage management, such as data locations, storage management components associated with data, storage devices used in performing storage operations, index data, data application type, or other data.

With this arrangement, if it is desired to obtain information regarding network data, a system administrator or system process may simply consult the metabase for such information rather than iteratively access and analyze each data item in the network. Thus, this significantly reduces the amount of time required to obtain data object information by substantially eliminating the need to obtain information from the source data, and furthermore minimizes the involvement of network resources in this process, substantially reducing the processing burden on the host system.

Various embodiments of the invention will now be described. The following description provides specific details for a thorough understanding and enabling description of these embodiments. One skilled in the art will understand, however, that the invention may be practiced without many of these details. Additionally, some well-known structures or functions may not be shown or described in detail, so as to avoid unnecessarily obscuring the relevant description of the various embodiments.

The terminology used in the description presented below is intended to be interpreted in its broadest reasonable manner, even though it is being used in conjunction with a detailed description of certain specific embodiments of the invention. Certain terms may even be emphasized below; however, any terminology intended to be interpreted in any restricted manner will be overtly and specifically defined as such in this Detailed Description section.

A flow chart 100 illustrating some of the steps associated with one embodiment of the present invention is shown in FIG. 1. In order to perform some of the functions described herein, it may be necessary at the outset to install certain data classification software or data classification agents on computing devices within at least parts of the network (step 102). This may be done, for example, by installing classification software on client computers and/or servers within a given network. In some embodiments, classification agents may be installed globally on a computing device or with respect to certain subsystems on a computing device. The classification software may monitor certain information regarding data objects generated by the computers and classify this information for use as further described herein.

Next, at step 104, a monitor agent may be initialized. Such a monitoring agent may be resident or installed on each computing device similar to the deployment of classification agents described above and may be configured to monitor and record certain data interactions within each machine or process. For example, the monitor agent may include a filter driver program and may be deployed on an input/output port or data stack and operate in conjunction with a file management program to record interactions with computing device data. This may involve creating a data structure such as a record or journal of each interaction. The records may be stored in a journal data structure and may chronicle data interactions on an interaction by interaction basis. The journal may include information regarding the type of interaction that has occurred along with certain relevant properties of the data involved in the interaction. One example of such a monitor program may include Microsoft's Change Journal or similar program.

Prior to populating a metabase with metadata, the portions of the network or subject system may be quiesced such that no data interactions are permitted prior to completing an optional scan of system files as described in conjunction with step 106 below. This may be done in order to obtain an accurate point in time picture of the data being scanned and to maintain referential integrity within the system. For example, if the system were not quiesced, data interactions would continue and be allowed to pass through to mass storage and thus data would change. However, in some embodiments, the subject system may be allowed to continue to operate, with operations or instructions queued in a cache. These operations are typically performed after the scan is complete so that any such data interactions occur based on the cached operations are captured by the monitor agent.

Generally, the file scanning of step 106 may be performed by a data classification agent and may include traversing the file system of a client to identify data objects or other files, email or other information currently stored or present in the system and obtain certain information regarding the information such as any available metadata. Such metadata may include information about data objects or characteristics associated with data objects such as the data owner (e.g., the client or user that generates the data or other data manager), the last modified time (e.g., the time of the most recent modification), the data size (e.g., number of bytes of data), information about the data content (e.g., the application that generated the data, the user that generated the data, etc.), to/from information for email (e.g., an email sender, recipient or individual or group on an email distribution list), creation date (e.g., the date on which the data object was created), file type (e.g., format or application type), last accessed time (e.g., the time the data object was most recently accessed or viewed), application type (e.g., the application which generated the data object), location/network (e.g., a current, past or future location of the data object and network pathways to/from the data object), frequency of change (e.g., a period in which the data object is modified), business unit (e.g., a group or department that generates, manages or is otherwise associated with the data object), and aging information (e.g., a schedule, which may include a time period, in which the data object is migrated to secondary or long term storage), etc. The information obtained in the scanning process may be used to initially populate the metabase of information regarding network data at step 108.

After the metabase has been populated, the network or subject system may be released from the quiesced state and normal operation may resume. Next, at step 110, the monitor agent may monitor system operations to record changes to system data in the change journal database as described above. The change journal database may include a database of metadata or data changes and may comprise log files of the data or metadata changes. In some embodiments, the data classification agent may periodically consult the change journal database for new entries. If new entries exist, these entries may be examined, and if deemed relevant, then analyzed, parsed, and written to the metabase as an update (step 112). In other embodiments, change journal entries may be supplied substantially in parallel to the journal database and data classification agent. This allows the metabase to maintain substantially current information regarding the state of system data at any given point in time.

As mentioned above, one benefit of such a metabase is it significantly reduces the amount of time required to obtain information by substantially eliminating the need to obtain information directly from the source data. For example, assume a system administrator desires to identify data objects

that a certain user has interacted with that contain certain content or other characteristics. Rather than search each file in each directory, a very time consuming process, the administrator may simply search the metabase to identify such data objects and any properties associated with those objects, (e.g., metadata, which may include location, size, etc.), resulting in significant time savings.

Moreover, use of the metabase for satisfying data queries also minimizes the involvement of network resources in this process, substantially reducing any processing burden on the host system. For example, as described above, if an administrator desires to identify certain data objects, querying the metabase rather than the file system virtually removes the host system from the query process (i.e., no brute force scanning of directories and files required), allowing host computing devices to continue performing host tasks rather than be occupied with search tasks.

FIG. 2 shows one embodiment of a client 200 constructed in accordance with principles of the present invention. As shown, client 200 may include a classification agent 202 and a monitor agent 206, which, in some embodiments, may be combined as an update agent 204, and which may be a single module encompassing the functionality of both agents. Client 200 may also include an internal or external data store 209, metabase 210, and change record 212.

Generally, client 200 may be a computing device, or any portion of a computing device that generates electronic data. Data store 209 generally represents application data such as production volume data used by client 200. Metabase 210, which may be internal or external to client 200 may contain information generated by classification agent 202 as described above. Similarly, change journal 212, which also may be internal or external to client 200, may contain information generated by monitor agent 206 as described above.

In operation, data interactions occurring within client 200 may be monitored with update agent 204 or monitor agent 206. Any relevant interaction may be recorded and written to change record 206. Data classification agent 202 may scan or receive entries from monitor agent 206 and update metabase 210 accordingly. Moreover, in the case where update agent 204 is present, monitored data interactions may be processed in parallel with updates to change record 212 and written to data store 208 and metabase 210 occurring accordingly. A file system 207 may be used to conduct or process data from the client to a data store 209.

FIG. 3a provides a more detailed view of the journaling and classification mechanisms of client 200 generally shown in FIG. 2. As shown, system 300 may include a memory 302, an update agent 304 which may include a separate or integrated monitor agent 306, classification agents 312a and 312b, a content agent 315, a monitor program index 310, metabase 314 and mass storage device 318.

In operation, data interactions that occur between memory 302 and mass storage device 318 may be monitored by monitor agent 306. In some embodiments, memory 302 may include random access memory (RAM) or other memory device employed by a computer in client 200 in performing data processing tasks. Certain information from memory 302 may be periodically read or written to mass storage device 318 which may include a magnetic or optical disk drive such as a hard drive or other storage device known in the art. Such data interactions are monitored by monitoring agent 306 which, in some embodiments, may include any suitable monitoring or journaling agent as further described herein.

As shown, system 300 may also include an administrative file system program 316, such as a file system program known in the art, which may include operating system programs, a

FAT, an NTFS, etc. that may be used to manage data movement to and from mass storage device 318. Thus, in operation, data may be written from memory 302 to mass storage device 318 via file system program 316. Such an operation may occur, for example, to access data used to service an application running on a computing device. During this operation, monitor agent 306 may capture this interaction and generate a record indicating that an interaction has occurred and store the record in index 310. The data may be stored in mass storage 318 under the supervision of file system manager 316.

As shown in FIG. 3a, monitor agent 306 may analyze data interactions such as interactions between memory 302 and mass storage 318 via the file system manager 316, and record any such interactions in monitor index 310. Thus, as described above, monitor index 310 may represent a list of data interactions wherein each entry represents a change that has occurred to client data along with certain information regarding the interaction. In embodiments where Microsoft Change Journal or other similar software is used, such entries may include a unique identifier such as an update sequence number (USN), certain change journal reason codes identifying information associated with a reason(s) for the change made, along with data or metadata describing the data and certain data properties, data copy types, etc.

Thus, in operation, as data moves from memory 302 to mass storage 318 (or vice versa), monitor agent 304 may create and write an entry to index 310 which may in turn, be analyzed and classified by classification agent 312b for entry in metabase 314. In some embodiments, classification agent 312a may be coupled with mass storage device (either directly or through file system manager 316) and write metadata entries to both metabase 314 and mass storage device 318. In some embodiments, the metabase information may be stored on mass storage device 318. Moreover, in an alternate embodiment, classification agent 312b may periodically copy or backup metabase 314 to the storage device under the direction of a storage manager and/or pursuant to a storage policy (not shown) such that the information in metabase 314 may be quickly restored if lost, deleted or is otherwise unavailable.

In some embodiments, optional classification agent 312a may operate in conjunction with monitor agent 306 such that data moving to mass storage device 318 is classified as further described herein and written to device 318. With this arrangement, the data, along with the processed metadata describing that data, is written to mass storage device 318. This may occur, for example in embodiments in which monitor agent 306 and classification agent 312a are combined into update agent 304. Writing metadata in this way allows it to be recovered or accessed from mass storage device 318 if necessary, for example, when metabase 314 is missing certain information, busy, or otherwise inaccessible.

Content agent 315 may be generally used to obtain or filter data relating to content of the data moving from memory 302 to mass storage 318. For example, content agent 315 may read data payload information and generate metadata based on the operation for storage in metabase 314 and may include a pointer to the data item in mass storage 318. The pointer information may optionally be stored in an index. This metadata may also be stored with the data item in mass storage 318 or as an entry functioning in place of or in addition to metabase 314. Storing metadata relating to data content in metabase 314 provides the capability to perform content searches for data in the metabase 314, instead of searching entries in mass storage 318. This allows the system to quickly locate information satisfying a content query that may be retrieved from mass storage 318, if necessary.

Moreover, such content metadata may be generated and used in locating data based on content features throughout a hierarchy within a storage operation system (e.g., content metadata may be generated and stored at each or certain levels of storage within the system (primary, secondary, tertiary, etc.) to facilitate location and retrieval of data based on content). As will be understood by one of skill in the art, the functionality provided by the content agent 315, classification agent 312a & b and monitor agent 306 may be provided by one or more modules or components such that the modules may be integrated into a single module providing the described functions, or may be implemented in one more separate modules each providing some of the functions.

FIG. 3b is a flow chart 350 illustrating some of the steps that may be involved in the journaling process described above. At step 355, the monitor program may be initialized, which may include instantiating a data structure or index for recording interaction entries, and the assignment of a unique journal ID number which allows the system to differentiate between various journaling data structures that may be operating within the system. As mentioned above, the monitor program may include a filter driver or other application that monitors data operations (step 360). During the monitoring process, the monitor agent may observe data interactions between memory and mass storage to determine that certain data interactions have occurred. Information relating to the interactions may be captured and used to populate the metabase. In some instances, interaction types or certain aspects of interactions are captured. Such types or aspects may be defined in an interaction definition, which may be a change journal reason codes as used by Microsoft's Change Journal program, or be defined by a user or network administrator to capture some or all data interactions in order to meet user needs. For example, certain change definitions may record every data interaction that occurs regardless of whether any data actually changes or not. Such information may be useful, for example, to determine users or processes that have "touched scanned or otherwise accessed data without actually changing it.

Thus, it is possible to employ interaction definitions that may capture a relatively broad or narrow set of operations, allowing a user to customize the monitor program to meet certain goals. Such interaction definitions may define or describe data movement, changes, manipulations or other operations or interactions that may be of interest to a system user or administrator (e.g., any operation that "touches" data may be recording along with the action or operation that caused the interaction (e.g. read, write, copy, parse, etc.) Moreover, change definitions may evolve over time or may be dynamic based on the entries sent to the index. For example, if expected results are not obtained, change definitions may be modified or additional definitions used until appropriate or desired results are obtained. This may be accomplished, for example by globally linking certain libraries of interaction definitions and selectively enabling libraries on a rolling basis until acceptable results are achieved. This process may be performed after the initial activation of the monitor agent and periodically thereafter, depending on changing needs or objectives.

Moreover, in some embodiments, the system may support the use of "user tags" that allow certain types of information to be tagged so they may be identified and tracked throughout the system. For example, a user may designate a particular type of data or information such as project information, or information shared between or accessed by particular group of users to be tracked across the system or through various levels of storage. This may be accomplished through a user

interface (not shown) that allows a user to define certain information to be tagged, for example, by using any available attribute within the system such as those specified above with respect to the classification agent or filter used in the system. In some embodiments, the user may define one or more tags using these or other attributes which may be further refined by combining them through the use of logical or Boolean operators to a define a certain tag expression.

For example, a user may define a certain tag by specifying one or more criteria to be satisfied such as certain system users, a certain data permission level, a certain project, etc. These criteria may be defined using a logical operators such as an AND or OR operators to conditionally combine various attributes to create a condition that defines a tag. All information satisfying those criteria may be tagged and tracked within the system. For example, as data passes through the monitor agent 306 (or other module within update agent 304), the data satisfying these criteria may be identified and tagged with a header or a flag or other identifying information as is known in the art. This information may be copied or otherwise noted by metabase 314 and mass storage 318 so that the information may be quickly identified. For example, the metabase may contain entries keeping track of all entries satisfying the tag criteria along with information relating to the types of operations performed on the information as well as certain metadata relating to the data content and its location in mass storage 318. This allows the system to search the metabase at a particular level of storage for the information, and quickly locate it within mass storage device for potential retrieval.

Next, a step 365, the monitor agent may continue to monitor data interactions based on change definitions until an interaction satisfying a definition occurs. Thus, a system according to one embodiment of the present invention may continue to monitor data interactions at steps 360 and 365 until a defined interaction, such as an interaction satisfying or corresponding to a selection criterion, such as an interaction definition, etc., occurs. If a defined interaction does occur, the monitor agent may create a record, which may be stored in a monitor program index, and in some embodiments, assign an interaction code that describes the interaction observed on the data object. Next, a step 370, the monitor program may identify a data object identifier associated with the data and that is associated with the data interaction, such as a file reference number (FRN) related to the data object. The FRN may include certain information such as the location or path of the associated data object. Any additional information (e.g., data properties, copy properties, storage policy information, etc.) associated with the FRN may also be obtained in order to enrich or enhance the record. In some embodiments, this may further involve obtaining information from other system files including master file tables (MFTs) to further enhance the metabase entries. Additional processing or formatting of the metabase entries may also occur in accordance with certain defined classification paradigms in order to populate the metabase with optimal or preferred information.

Next, at step 375 the record may be assigned record identifier such as a unique update sequence number (USN) that may be used to uniquely identify the entry within the index, and, in some embodiments, act as an index memory location. Thus a particular record may be quickly located with a data structure based on the USN. Next at step 380, the information described above may be concatenated or otherwise combined with other data or metadata data obtained by the monitor agent and arranged in an appropriate or expected format to create the record that may be written to the monitor index.

In alternate embodiments, the information described above may be written to the index and arranged at the index into an

expected format or may be written to the record "as received" and include metadata or other information, such as a header describing that information such that adherence to a strict data format is not required. For example, some records may contain more or less information than other records, as appropriate. After the record has been constructed and deemed complete, the record may be "closed" by the system at step 385 and the system may then assign another USN to the next detected change. However, if the record is determined to be incomplete, the monitor agent or update agent may request any missing information to complete the entry. If such information is not received, the monitor agent may set a flag within the record to indicate it contains incomplete information and the record may then be closed.

FIG. 4 is a flow chart 400 illustrating some of the steps that may be involved in a data classification process. At step 410, the classification agent may be initialized, which may include activating, reserving and/or clearing certain buffers and/or linking to libraries associated with deployment of the classification agent. Prior to scanning the interaction records generated by the monitor agent as described above, the classification agent may classify existing stored data by, for example, traversing the file and directory structure of an object system to initially populate the metabase as described herein.

Next, at step 420, during normal operation, the classification agent may scan the entries in the interaction journal to determine whether new entries exist since any previous classification processing was performed, for example, by determining whether the most recent entry currently in the journal is more or less recent than the last journal entry analyzed. This may be accomplished in several ways. One method may include scanning a time or date information associated with the last journal entry examined and comparing it to the most recent time or date information than the entry currently present in the journal. If it is determined that the most recent journal entry occurred after a previous classification process, this process may be performed iteratively by "walking backwards" through the journal entries until the last entry previously analyzed by the classification agent is found. All entries with time information after that point may be considered new or unprocessed by the classification agent (step 440). If the last entry analyzed has the same time stamp as the most recent entry in the journal, no new entries exist and the system may return to step 420 to continue monitoring, etc.

Another method of identifying new journal entries may include comparing record identifiers such as USN numbers assigned to each journal entry (step 430). Journal entries with a larger USN number than the last entry previously analyzed may be considered new or unprocessed. If the last entry analyzed has the same USN number as the current entry, no new entries exist and the system may return to step 420 to continue monitoring, etc. This comparison may be performed until new entries are located (step 440) or until it is determined that no new entries exist.

In other embodiments, rather than scanning the journal data structure for new entries, any entries created by the journal agent may be automatically sent to the classification agent and the identification process may be largely unnecessary (except for the case where such scanning is necessary or desirable, for example, to repopulate the journal or verify certain information, etc.).

Next, at step 450, assuming new journal entries are identified, the system may determine if a metabase record already exists for the data object associated with those entries. This may be accomplished by comparing data object identifiers, such as FRNs of metabase entries with data object identifiers

such as FRNs of journal entries. Matching these and other unique data characteristics may be used to match or correlate metabase and journal entries.

If no corresponding metabase record is identified, a new record may be created at step 460. This may involve creating a new metabase entry ID, analyzing the journal entry and parsing the entry into a predetermined format, and copying certain portions of the parsed data to the new metabase entry (steps 460 and 470), as further described herein. Any additional metadata or file system information may also be associated with the new entry to enhance its content, including information from an FRN or information derived from an interaction code present in the journal entry, file system, e.g., MFT, etc. (step 480).

On the other hand, if a corresponding metadata entry is identified, the new journal entry may be processed as described above and may overwrite some or all of the corresponding entry. Such an updated pre-existing entry may receive an updated time stamp to indicate a current revision. However, in some embodiments, even if a corresponding entry is located, a new entry may be created and written to the metabase and optionally associated with the existing record. In such a system, the older related record may be maintained, for example, archival, recreation, historical or diagnostic purposes, and in some embodiments, may be marked or indicated as outdated or otherwise superseded. Such corresponding entries may be linked to one another via a pointer or other mechanism such that entries relating to the history of a particular data object may be quickly obtained.

Next, at step 490 the system may process any additional new journal entries detected by returning to step 450, where those additional new entries may be processed as described above. If no new entries are detected, the system may return to step 420 to perform additional scans on the journal data structure and continue monitoring.

FIG. 5 illustrates an embodiment of the present invention in which a secondary processor performs some or all of functions associated with the data classification process described herein, including certain search functions. As shown, system 500 may include a manager module 505 which may include an index 510, a first computing device 515, (which may include a first processor 520, a journal agent 530, and a data classification agent 535), and a second computing device 540 which may include a second processor 545 and a data classification agent 535. System 500 may also include data store 550, a metabase 555 and change journal 560.

Computing devices 515 and 544 may be any suitable computing device as described herein and may include clients, servers or other network computers running software, such as applications or programs that create, transfer, and store electronic data. In some embodiments, metabase 555 and journal 560 may be physically located within computing device 515, e.g., stored on local mass storage. In other embodiments the metabase 555 and journal 560 may be external to computing device 515 (or distributed between the two). In yet other embodiments, metabase 555 is accessible via a network and journal 560 is a local device.

In operation, computing device 515 may operate in a substantially similar manner as system 300 shown in FIG. 3 with second processor 545 in second computing device 540 performing certain functions. For example, as shown, data classification agent 535 and journaling agent 530 may operate substantially as described in connection with FIG. 3, i.e., journaling agent monitors data interactions on computing device 515 and records the interactions in journal 535 and classification agent processes journal entries and populates metabase 555, etc.

However, certain of the functions may be initiated or performed in whole or in part by second processor 545. Computing operations associated with journal agent 530 and/or classification agent 535 may run on or be directed by second processor 545 and may also utilize support resources located on or associated with computing device 540 such that the resources on computing device 515 are substantially unimpacted by these operations. This may serve to offload certain non-critical tasks from the host system (515) and have them performed by a secondary computing device 545.

For example, in some embodiments, the processing burden associated with some or all of the following tasks normally performed by first computing device 515 may be performed by processor 545 and associated resources in second computing device 540: (1) the initial scan of client files by the classification agent 535 and population of metabase 555, (2) the ongoing monitoring of data interactions of computing device (e.g., 515) and generation of interaction records for storage in journal 560, (3) processing and classification of journal information for updating metabase 555; and (4) searching or otherwise analyzing or accessing metabase 555 and/or journal 560 for certain information. However, in some embodiments it may be preferred to assign the secondary computing device the certain tasks such as those associated with searching metabase 555, while other tasks such as updating the journal and metabase may be performed by the primary computing device.

Performing such operations using a secondary or other processor may be desirable, for example, when the primary processor (e.g., processor 520) is unavailable, over utilized, unavailable or otherwise heavily used, or when it is otherwise desired to remove the primary processor and other primary system resources from performing certain tasks such as the ones described above. For example, in the case where it is desired to search or access metabase 555 for certain information, it may be preferable to avoid using processor 520 for this task so it remains free to perform other tasks associated with programs operating on computing device 515 (e.g., when computing device 515 is busy performing other network or application-related functions).

In some embodiments, the secondary processor may be located on computing device 515 (e.g., processor 525) and may perform the operations described herein in connection with processor 545. Moreover, some embodiments may include a manager module 505 which may coordinate overall operations between the various computing devices. For example, manager module 505 may monitor or otherwise be cognizant of the processing load on each computing device and may assign processing tasks based on availability (e.g., load balance). For example, if processor 520 is idle or operating at a low capacity, a request to search metabase 555 may be handled by processor 520. However, if processor 520 is busy or scheduled to perform or is performing priority work, manager 505 may assign the task to processor 545. Manager 505 may act as a general arbiter for such processor assignments to ensure system 500 is making efficient use of system resources.

FIG. 6 is flow chart 600 illustrating some of the steps involved in performing a query on a metabase in a multiple processor system similar to the system shown in FIG. 5. At step 610, a query may be received by the system for certain information. This request may be processed and analyzed by a manager module or other system process that determines or otherwise identifies which metabase or metabases within the system likely include at least some of the requested information, step 630. For example, the query itself may suggest which metabases to search and/or the management module

may consult an index that contains information regarding metabase content within the system as further described herein. It will be understood that the identification process may require searching and identifying multiple computing devices within an enterprise or network that may contain information satisfying search criteria.

In other embodiments, search requests may be automatically referred to a secondary processor to minimize processing demands on the computing device that may have created or is otherwise associated with the identified metabase(s). In some embodiments, it is preferable that the computing device that created or is otherwise associated with the identified metabase(s) not be involved in processing search operations as further described herein. Thus, the secondary computing device may consult with a manager or index associated with other computing devices to identify metabases with responsive information.

Next at step 640, the secondary processor may search metabases to identify appropriate data set that may potentially have information related to the query. This may involve performing iterative searches that examine results generated by previous searches and subsequently searching additional, previously unidentified metabases to find responsive information that may not have been found during the initial search. Thus the initial metabase search may serve as a starting point for searching tasks that may be expanded based on returned or collected results. Next, at step 650, the returned results may be optionally analyzed for relevance, arranged, and placed in a format suitable for subsequent use (e.g., with another application), or suitable for viewing by a user and reported (step 650).

FIG. 7 presents a system 700 constructed in accordance with the principles of the present invention employing a centralized metabase 760 that may serve multiple computing devices 715-725. For example, as shown, system 700 may include computing devices 715-725, each of which may include a journaling agent (730-740 respectively), a classification agent (745-755 respectively), and centralized metabase 760, and in some embodiments, a manager module 705 with an index 710.

In operation, system 700 may operate substantially similarly to system 300 shown in FIG. 3 with each computing device 715-725 storing classification entries in centralized metabase 760 rather than each computing device having its own dedicated metabase. For example, as shown, data classification agents 745-755 may operate substantially as described herein and communicate results to centralized metabase 760. That is, analyze and process entries within the respective journals associated with journaling agents 730-740, and report results to metabase 760. With this arrangement, the classification agent may provide each metabase entry with an ID tag or other indicia that identifies which computing device 715-725 the entry originated from to facilitate future searches and efficiently designate entry ownership, or other associations between entries and computing devices.

Moreover, each entry to metabase 760 may be assigned a unique identifier for management purposes. As mentioned above, this number may represent the index location or offset of the entry within centralized metabase 760. In some embodiments, entries may be communicated to metabase 760 from the computing devices 715-725 on a rolling basis and may be arranged and formatted for storage by the metabase 760. For example, metabase 760 may receive multiple entries at substantially the same point in time from multiple computing devices 715-725 and may be responsible for queuing and arranging such entries for storage within the metabase 760.

In some embodiments, system 700 may include manager module 705 that may be responsible for assigning or removing associations between certain computing devices 715-725 and a particular centralized metabase 760. For example, in accordance with certain system preferences defined in index 710, manager 705 may direct certain computing devices 715-725 to write classification entries to a particular centralized metabase 760. Information indicating an association of the metabase 760 and the computing devices 715-725 may be stored in the index 710. This allows system 700 to reassign resources (globally or locally) to optimize system performance without the need to change device pointers or code associated with each computing device 715-725 that may be affected by a particular reallocation. For example, manager 705 may reassign certain computing devices 715-725 to another metabase by changing a destination address in an appropriate index.

FIG. 8 is flow chart 800 illustrating some of the steps involved in using a centralized metabase with multiple computing devices similar to the one shown in FIG. 7. At step 810, a centralized metabase may be instantiated by a manager module or in accordance with certain system management or provisioning policies. This may involve securing certain processing, storage, and management resources for performing the task, loading certain routines into various memory buffers and informing the management module that the metabase is ready for operation.

Next, at step 820, the management module may review system resources, management policies, operating trends, and other information, for example, to identify computing devices to associate with the instantiated centralized metabase. This may further involve identifying pathways to the metabase from the various computing devices, locating operational policies governing the computing devices and, creating certain logical associations between the centralized metabases and the identified computing devices. These associations, once created, may be stored in an index or database for system management purposes.

After the metabase has been instantiated and associated with computing devices, classification agents within each associated computing device may scan existing files or data on the computing devices or clients (step 825) and populate the centralized metabase as further described herein (step 830). During the scanning process, a computing device identifier or other indicia may be appended or otherwise associated with the entry prior to transmission to the metabase such that each entry in the metabase can be tracked to its associated source computing device (step 840). Next, the centralized metabase may be populated with entries (step 850) and may communicate with the management module to establish and monitor a list of computing devices serviced by the centralized metabase and return to step 830. At this point, the system continues to monitor the associated computing devices for data interactions, which may be reported to the centralized metabase on an ongoing, periodic, or rolling basis.

In certain circumstances, the centralized metabase may need to assimilate or otherwise integrate existing entries with new entries reported by the computing devices. For example, the centralized metabase may become disconnected or unavailable for a period of time and subsequently be required to integrate a large number of queued entries. In this case, the metabase or management module may examine existing metabase entries as described herein and communicate with computing devices to identify: (1) the amount of time the object computer and the metabases have been disconnected, (2) the number of queued entries at the computing devices that need to be processed (for example, entries cached once

the centralized metabase was inaccessible for write operations), (3) whether there are any duplicative entries, and (4) which entries need to be integrated and in what order of preference (assuming multiple computing devices contain queued entries).

Based on these criteria, the management module or centralized metabase may assimilate the relevant entries into the metabase in the appropriate order until the backlog is eliminated and the system returns to normal operation. If it is determined during this process that certain information is lost to cache overflow, accidental deletion, corruption, or other reasons, the metabase and/or manager module may indicate such a discontinuity with the metadata or index associated with the centralized storage device or management module. In this case, clients, computing devices or other data sources may be rescanned to replace or repair the faulty entries. In other embodiments, the points of discontinuity may be noted and interpolation or other data healing techniques may be employed to provide derived information for the unknown points based on known information.

FIG. 9 presents a system 900 constructed in accordance with the principles of the present invention including a computing device that interacts with a network attached storage device (NAS). As shown, system 900 may include a management module 905 and index 910, computing devices 915-925, each of which may include a journaling agent (945-955 respectively), a classification agent (930-940 respectively), data stores 960 and 965, and metabases 970-980. System 900 may also include NAS device 995 which may include NAS storage device 990 and NAS file system manager 985. Moreover, computing device 925 may be configured to operate as a NAS proxy device supervising the transfer of data to and from NAS device 995.

In operation, system 900 may operate substantially similar to system 300 shown in FIG. 3a with exception of the NAS portion shown on the right-hand side. For example, as shown, data classification agents 930-940 may operate substantially as described herein and communicate results to their respective metabases 970-980. That is, analyze and process entries within the respective journals associated with journaling agents 945-955, and report results to metabases 970-980 which may be supervised in whole or in part by management module 905.

Data from computing device 925 may be journaled and classified using methods similar to those described herein. For example, journaling agent 955 may reside on computing device 925 and track each or certain data interactions between NAS device 995 and external applications. The location of the journaling agent 955 may be external to the NAS device 995 due, at least in part, to its proprietary nature (i.e., a closed system) and the difficulty associated with attempting to run other programs on the NAS device 995 itself.

The NAS portion 995 of system 900 may operate somewhat differently. For example computing device 925 may operate as a NAS proxy for moving data files to and from NAS device 995 using a specialized protocol such as the Network Data Management Protocol (NDMP) that is an open network protocol designed to perform data backups over heterogeneous networks. NDMP may be used to enhance performance by transferring data over a network by separating data and control paths, while maintaining centralized backup administration.

Journaling agent 955 may record any interactions between NAS data and external applications and record those interactions in computing device 925 as described herein. In some embodiments, such a journaling agent may include specialized routines for interpreting and processing data in NAS

format. Data classification agent 940 may analyze journal entries and populate metabase 980 initially and periodically as further described herein.

Once initially populated, it may be desired to search the metabases of system 900 for certain information. This is discussed in more detail below in connection with the flow chart 1100 of FIG. 11. In some embodiments, this may be handled by manager 905 or other system process which may initially evaluate any search request and consult index 910 or other information stores to determine which metabases within the system are likely to include responsive information. The results of this evaluation may be provided to the computing device handling the search request and may be in the form of pointers or other indicia or identifiers identifying a metabase such as a metabase ID. This may allow the computing device posing the search request to contact and search the identified metadata directly. In other embodiments manager 905 may process the request and provide substantially complete results to the computing device that submitted the query.

FIG. 10 is flow chart 1000 illustrating some of the steps that may be involved in using the NAS system similar to or the same as the one shown of FIG. 9. At step 1010 a copy operation may be initiated that directs data from computing device to a NAS device. This may involve identifying certain data to moved, for example, based on a data management or storage policy. Other factors that may also be considered may include data size, the last time the data was moved to the NAS device, the file owner, application type, etc.

It will be understood that in some embodiments it may be preferred to use computing device 925 as a NAS proxy that routes data from other network computing devices (not shown) to NAS device 995 with the computing device 925 supervising the data movement using certain specialized transfer programs to assist in the effort (step 1020). As the data is routed through computing device 925, journaling agent 955 may monitor interactions with NAS device 995 and create interaction entries for an interaction journal (step 1030). This may be accomplished by consulting with NAS file manager 985 and identifying which files in NAS 995 that have been involved in a data interaction as further described herein (step 1040). Next, journal entries may be created or updated to reflect data interactions currently detected as previously described herein (step 1050). The interaction journal may then be scanned to analyze the journal records (step 1060) and perform the classification process as further described herein to create metabase entries (step 1070). At this point metabase entries may be assigned an identifier and used to populate metabase 980 (step 1080).

As mentioned above, under certain circumstances, it may be desired to search a system that includes multiple metabases for certain information such as system 900 shown in FIG. 9 whether or not NAS included. FIG. 11 includes a flow chart 1100 illustrating some of the steps that may be performed in searching a multiple metabase system in accordance with certain aspects of the present invention.

Assume, for example, a user wants to locate and copy all data relating to a certain specified criteria such as data relating to a specific marketing project created and edited by a specific group of users over a certain period of time. First, the requestor may formulate such a request through a user interface (not shown) using techniques known in the art and submit the request to the system for processing. This may also be accomplished by an automated computerized process, for example, when the system is performing certain management functions. Next the system may receive and analyze this query (step 1110). In some embodiments, this may be per-

formed by a computing device configured to support the user interface. In other embodiments, the computing device may simply pass the request to the system where a management module or other system process computing device may perform the analysis. The analysis may include determining characteristics of data in the metabase that may satisfy the selected criteria.

Once the search request or query has been analyzed or parsed, the system may identify all metabases likely to contain records related to relevant data objects based on a query. This may be accomplished by using information obtained from analyzing or parsing the request as well as consulting with a management module that may have a substantially global view of metabases within the system that includes index information or a general overview of the information the metabases contain. After a set of metabases have been identified, the management module or other computing device may perform the search to identify a data set satisfying a query as further described herein and return a set of results (step 1130). At step 1140 the results may optionally be normalized. If normalization is not required, the results may be reported at step 1150. If normalization is desired, the system may analyze the results for both content and completeness. If, based on the returned results, other unsearched metabases are implicated as potentially having information that satisfies the search criteria, those metabases may be searched as well. This process may continue in an iterative fashion until a substantially complete set of results is obtained. Even if no additional metabases are implicated, these results may then be optionally normalized by performing certain functions such as locating and removing duplicative results, identifying network pathways to data objects identified in the search, and formatting or arranging the results for further processing (whether another computing process or for a user). For example, the returned results may be used to locate and retrieve the responsive data objects that may include information located on primary or secondary storage devices within the system or for other purposes as further described herein.

In some embodiments, the systems and methods of the present invention may be employed to identify and track some or all data interactions on a user or group basis. For example, a system administrator or user may wish to record and keep track of all data interactions involving some or all system groups or users. This may include, for example, read and write operations performed on the user's or group's behalf, information and applications used or accessed, viewed web pages, electronic gaming interactions, chat, instant messages, and other communication interactions, multimedia usage, and other Internet or network based electronic interactions as known in the art. Thus, the system identifies, captures, classifies, and may otherwise track user and group interactions with electronic data creating a data store or other repository of these interactions and metadata associated with these interactions. In some embodiments, this repository may serve as a "digital or electronic life record" that effectively chronicles and catalogues some or all user or group interactions with electronic information and data during a given time period as further described herein.

For example, FIG. 11a illustrates a system constructed in accordance with the principles of the present invention that identifies, captures, classifies, and otherwise tracks user and group interactions with electronic data. As shown, the system may generally include computing device 1162, one or more classification agents 1164, one or more journaling agents 1165, metabase 1166, change record 1167, and database 1168.

In operation computing device 1162 may be coupled to or interact with various other applications, networks, and electronic information such as, for example multimedia applications 1170, instant messaging/chat applications 1172, network applications 1174 such as an enterprise WAN or LAN, Internet 1176, and gaming applications 1178. It will be understood, however, that these are only examples and that any other network, application, or type of electronic information suitable for the purposes described herein may be added if desired.

Journaling agents 1165 and classification agents 1164 may operate in conjunction with one another to detect and record data interactions as further described herein. For example, each type of electronic data interaction (e.g.—instant messaging, web surfing, Internet search activities, electronic gaming, multimedia usage, etc.) may be identified, captured, classified, and otherwise tracked by a different journaling agent 1165 and classification agent 1164, for example an interaction-specific journaling agent 1165 or classification agent 1164 dedicated to processing a single type of interaction with electronic data. Thus, the system may have a first journaling agent 1165 and a first classification agent 1164 monitoring network traffic on a given network interface (not shown) directed to interactions associated with Internet usage, and a second journaling agent 1165 and a second classification agent 1164 monitoring a different system resource directed to interactions associated with electronic gaming (e.g.—recording and classifying gaming interactions such as recording games played, opponents played, win/loss records, etc.) or directed to interactions associated with use of an Internet browser to "surf" web (e.g.—tracking pages visited, content, use patterns, etc.) In some embodiments, journaling agent 1165 and classification agent 1164 may function as a single module capable of performing some or all functions associated with journaling agent 1165 and a classification agent 1164.

Thus, as a user or group interacts with various types of electronic information, some or all of those interactions may be captured and recorded in database 1168. Change record 1167 and metabase 1166 may record certain aspects of the interactions as further described herein and may represent an interaction by interaction log of the user's computing activities.

For example, in operation, a user of computing device 1162 may interact with certain applications such as multimedia application 1170 and instant messaging application 1172. This may include sending, receiving, viewing and responding to various audio/video files in any suitable format and may include instant, text or email messages. Journaling agent 1165 may detect the interactions between these applications and computing device 1162 and classification agent 1164 may classify and record information (e.g., metadata) associated with these interactions in metabase 1166 as further described herein.

Moreover, in some embodiments, some or all the content being exchanged or otherwise associated with these interactions may be captured and stored in database 1168 or other storage locations in the system. This may include capturing screen shots or summaries of information exchanges during data interactions. For example, the system may download all content associated with web pages viewed thus being able to recreate the original page content and interaction without access to the original or source version of the page on the Internet or other network. This may be advantageous, for example, if a user wishes to interact with content associated with a previous interaction when that content is no longer available, as is common with web pages and other network

21

resources over time. As another example, the system may also capture or otherwise store data associated with other interactions, for example chat transcripts, video game replays, search queries, search results, and associated search content, songs accessed, movies accessed, stored songs and movies, in addition to metadata, etc.

Moreover, in some embodiments, specialized classifications agents may be employed for some or all of the applications that a user or administrator desires to track and record. For example, the multimedia and instant messaging applications described above may each have a dedicated classification agent that analyzes journal records to create entries for metabase **1166**. Further still, each classification agent may have its own associated metabase and or repository for source data (not shown), so application histories and content may be quickly indexed and searched. In other embodiments, however, a “universal” classification agent may be used that recognizes the application type (e.g., based on the journaling agent entries) and process interactions accordingly (which may include routing metadata to one or more specialized metabases).

As shown in FIG. **11a**, computing device **1162** may also interact with various network applications **1174** such as LAN or WAN applications. These may include interaction with certain distributed programs such as Microsoft Word or Outlook. Users may also interact with Internet **1176** and download various web pages and other information. In accordance with an aspect of the present invention, interactions with these networks/applications may also be journaled as described above with certain information regarding these interactions stored in metabase **1166**. Portions of exchanged content may also be stored in database **1166**. For example, Word documents, emails, web pages, web addresses and HTML content may be captured and stored on database **1168** such that it substantially represents a record of all user interactions with computing device **1162**, or other system devices. For example, user interactions may be recorded with respect to any identified user based on identifiers and tracked at any network computing device.

Thus, if desired a user may retrieve captured data and review or replay certain data exchanges or save such records for future reference. For example, a user may store all instant messaging interactions for replay or transmission to another. In some instances, it may be desirable to not record certain interactions, such as personal or private information. In some embodiments, this may be accomplished by “disabling” the appropriate classification agent for a certain period of time, etc.

Likewise, interactions with gaming applications (network or stand alone) may also be recorded stored with appropriate information stored in database **1168** and metabase **1166**. Thus, a user may have the ability to retrieve, replay and transmit certain saved gaming sequences to third parties.

In some embodiments, database **1168** may become large and thus some information stored thereon may be moved to single instance storage from database **1168** with a pointer placed in the logical address of the instanced information (not shown). This may be performed as a memory saving measure as at least some of the entries in database **1168** are likely to be duplicative.

Some of the steps associated with the method generally described above are illustrated in chart **1200** of FIG. **12** and may include the following. At the outset, a group or user of interest may be identified based on certain user related information or other network characteristics (step **1210**). Such characteristics may include Active Directory privileges, network login, machine ID, or certain biometrics associated with

22

a user or group member. These characteristics may be combined together or associated with one another to create a user or group profile. Such profiles may be stored in a database or index within a management module of the system and may be used as classification definitions within the system. When it is desired to identify or classify data items associated with a particular interaction, the system may compare certain attributes of the data involved in a detected interaction and associate that interaction with a particular group or user based on profile information (step **1220**).

Such associations may be stored in a metabase created to keep track of user or group interactions. Thus, in one embodiment, the metabase essentially represents a list of all data interaction for a particular group or user. If desired, a list or copy of all the data items touched (e.g., interacted with) by a group or user may be quickly obtained.

In operation, the system may, through the use of a journaling agent or the like, monitor data interactions for a particular computing device as described herein. The interactions may be analyzed by a classification agent as described herein and associated with one or more profiles (step **1230**). The association may be recorded in an identified metabase(s) that keeps track of a user’s or group’s interactions (step **1240**) which may include references to the data object(s) identified, the attributes compared, and the basis for the association, etc. As discussed herein, the journaling agent may continue to monitor data interactions throughout operation, so that each metabase is updated and continues to accurately represent the data touched by a particular group or user. The identified metabases are associated with an identified group or user (step **1250**), such as by storing an indication of the association in an index.

FIG. **13** presents a system **1300** constructed in accordance with the principles of the present invention for communicating metadata and/or data objects between two or more computing devices. As shown, system **1300** may generally include first and second computing devices **1310** and **1320**, respectively, associated data stores **1330** and **1340**, and metabases **1350** and **1360**. Computing devices in system **1300** may store data objects and metadata in their respective metabases and data stores as further described herein. In certain situations, however, it may be desired to transfer certain metadata between metabases **1350** and **1360** and certain data objects between data stores **1330** and **1340**. This may be desirable for example, to move certain data from one computing device to another, to recreate a certain application at another location, or to copy or backup certain data objects and associated metadata.

FIG. **14** presents a flow chart **1400** illustrating some of the steps associated with moving data between the computing devices described above. First, at step **1410**, data objects and/or associated metadata may be identified for movement from one computing device to another. This may be accomplished by forming a query for certain data, such as a search for data that may be to be moved or copied pursuant to a data management or storage policy, or in response to a request to move data relating to certain processes or applications from one computing device to another, for any other suitable purpose such as disaster recovery, resource reallocation or reorganization, load balancing, etc.

At step **1420**, the query may be analyzed and a first data store associated with a first computer may be searched for data objects satisfying the search criteria. Data objects identified during this process may then be transferred to a second data store associated with a second computing device (step **1430**). Metadata associated with the transferred data objects may also be identified in a first metabase associated with the

first computing device and transferred to an appropriate second metabase associated with the second computing device (step 1440). Such a transfer may involve copying data objects and metadata from one data store and metabase to another, or in some embodiments, may involve migrating the data from its original location to a second location and leaving a pointer or other reference to the second location so the moved information may be quickly located from information present at the original location.

FIG. 15 illustrates a one arrangement of resources in a computing network in accordance with the principles of the present invention. As shown, storage operation cell 1550 may generally include a storage manager 1501, a data agent 1595, a media agent 1505, a storage device 1515, and, in some embodiments, may include certain other components such as a client 1585, a data or information store 1590, database 1511, jobs agent 1520, an interface module 1525, and a management agent 1530. Such system and elements thereof are exemplary of a modular storage system such as the CommVault QiNetix system, and also the CommVault GALAXY backup system, available from CommVault Systems, Inc. of Oceanport, N.J. and further described in U.S. patent application Ser. No. 09/610,738 which is incorporated herein by reference in its entirety.

A storage operation cell, such as cell 1550, may generally include combinations of hardware and software components associated with performing storage operations on electronic data. According to some embodiments of the invention, storage operations cell 50 may be related to backup cells and provide some or all of the functionality of backup cells as described in application Ser. No. 09/354,058, which is hereby incorporated by reference in its entirety. However, in certain embodiments, storage operation cells may also perform additional types of storage operations and other types of storage management functions that are not generally offered by backup cells.

In accordance with certain embodiments of the present invention, additional storage operations performed by storage operation cells may include creating, storing, retrieving, and migrating primary storage data (e.g., 1590) and secondary storage data (which may include, for example, snapshot copies, backup copies, HSM copies, archive copies, and other types of copies of electronic data) stored on storage devices 1515. In some embodiments, storage operation cells may also provide one or more integrated management consoles for users or system processes to interface with in order to perform certain storage operations on electronic data as further described herein. Such integrated management consoles may be displayed at a central control facility or several similar consoles distributed throughout multiple network locations to provide global or geographically specific network data storage information.

In some embodiments, storage operations may be performed according to various storage preferences, for example as expressed by a user preference or storage policy. A storage policy is generally a data structure or other information source that includes a set of preferences and other storage criteria associated with performing a storage operation. The preferences and storage criteria may include, but are not limited to, a storage location, relationships between system components, network pathway to utilize, retention policies, data characteristics, compression or encryption requirements, preferred system components to utilize in a storage operation, and other criteria relating to a storage operation. Thus, a storage policy may indicate that certain data is to be stored in a specific storage device, retained for a specified period of time before being aged to another tier of secondary

storage, copied to secondary storage using a specified number of streams, etc. A storage policy may be stored in a storage manager database 1511, to archive media as metadata for use in restore operations or other storage operations, or to other locations or components of the system.

A schedule policy may specify when to perform storage operations and how often and may also specify performing certain storage operations on sub-clients of data and how to treat those sub-clients. A sub-client may represent static or dynamic associations of portions of data of a volume and are typically mutually exclusive. Thus, a portion of data may be given a label and the association is stored as a static entity in an index, database or other storage location used by the system. Sub-clients may also be used as an effective administrative scheme of organizing data according to data type, department within the enterprise, storage preferences, etc.

For example, an administrator may find it preferable to separate e-mail data from financial data using two different sub-clients having different storage preferences, retention criteria, etc. Storage operation cells may contain not only physical devices, but also may represent logical concepts, organizations, and hierarchies. For example, a first storage operation cell 1550 may be configured to perform a first type of storage operations such as HSM operations, which may include backup or other types of data migration, and may include a variety of physical components including a storage manager 1501 (or management agent 1530), a media agent 1505, a client component 1585, and other components as described herein. A second storage operation cell may contain the same or similar physical components, however, it may be configured to perform a second type of storage operations such as SRM operations, and may include as monitoring a primary data copy or performing other known SRM operations.

Thus, as can be seen from the above, although the first and second storage operation cells are logically distinct entities configured to perform different management functions (i.e., HSM and SRM respectively), each cell may contain the same or similar physical devices in both storage operation cells. Alternatively, in other embodiments, different storage operation cells may contain some of the same physical devices and not others. For example, a storage operation cell 1550 configured to perform SRM tasks may contain a media agent 1505, client 1585, or other network device connected to a primary storage volume, while a storage operation cell 1550 configured to perform HSM tasks may instead include a media agent 1505, client 1585, or other network device connected to a secondary storage volume and not contain the elements or components associated with and including the primary storage volume. These two cells, however, may each include a different storage manager 1501 that coordinates storage operations via the same media agents 1505 and storage devices 1515. This "overlapping" configuration allows storage resources to be accessed by more than one storage manager 1501 such that multiple paths exist to each storage device 1515 facilitating failover, load balancing and promoting robust data access via alternative routes.

Alternatively, in some embodiments, the same storage manager 1501 may control two or more cells 1550 (whether or not each storage cell 1550 has its own dedicated storage manager 100). Moreover, in certain embodiments, the extent or type of overlap may be user-defined (through a control console (not shown)) or may be automatically configured to optimize data storage and/or retrieval.

Data agent 1595 may be a software module or part of a software module that is generally responsible for copying, archiving, migrating, and recovering data from client com-

puter **1585** stored in an information store **1590** or other memory location. Each client computer **1585** may have at least one data agent **1595** and the system can support multiple client computers **1585**. In some embodiments, data agents **1595** may be distributed between client **1585** and storage manager **1501** (and any other intermediate components (not shown)) or may be deployed from a remote location or its functions approximated by a remote process that performs some or all of the functions of data agent **1595**.

Embodiments of the present invention may employ multiple data agents **1595** each of which may backup, migrate, and recover data associated with a different application. For example, different individual data agents **1595** may be designed to handle Microsoft Exchange data, Lotus Notes data, Microsoft Windows 2000 file system data, Microsoft Active Directory Objects data, and other types of data known in the art. Other embodiments may employ one or more generic data agents **1595** that can handle and process multiple data types rather than using the specialized data agents described above.

If a client computer **1585** has two or more types of data, one data agent **1595** may be required for each data type to copy, archive, migrate, and restore the client computer **1585** data. For example, to backup, migrate, and restore all of the data on a Microsoft Exchange 2000 server, the client computer **1585** may use one Microsoft Exchange 2000 Mailbox data agent **1595** to backup the Exchange 2000 mailboxes, one Microsoft Exchange 2000 Database data agent **1595** to backup the Exchange 2000 databases, one Microsoft Exchange 2000 Public Folder data agent **1595** to backup the Exchange 2000 Public Folders, and one Microsoft Windows 2000 File System data agent **1595** to backup the client computer's **1585** file system. These data agents **1595** would be treated as four separate data agents **1595** by the system even though they reside on the same client computer **1585**.

Alternatively, other embodiments may use one or more generic data agents **1595**, each of which may be capable of handling two or more data types. For example, one generic data agent **1595** may be used to back up, migrate and restore Microsoft Exchange 2000 Mailbox data and Microsoft Exchange 2000 Database data while another generic data agent may handle Microsoft Exchange 2000 Public Folder data and Microsoft Windows 2000 File System data, etc.

Data agents **1595** may be responsible for arranging or packing data to be copied or migrated into a certain format such as an archive file which is discussed in more detail in connection with FIG. 16 herein. Nonetheless, it will be understood this represents only one example and any suitable packing or containerization technique or transfer methodology may be used if desired. Such an archive file may include a list of files or data objects copied in metadata, the file and data objects themselves. Moreover, any data moved by the data agents may be tracked within the system by updating indexes associated appropriate storage managers or media agents.

Generally speaking, storage manager **1501** may be a software module or other application that coordinates and controls storage operations performed by storage operation cell **1550**. Storage manager **1501** may communicate with some or all elements of storage operation cell **1550** including client computers **1585**, data agents **1595**, media agents **1505**, and storage devices **1515**, to initiate and manage system backups, migrations, and data recovery.

Storage manager **1501** may include a jobs agent **1520** that monitors the status of some or all storage operations previously performed, currently being performed, or scheduled to be performed by storage operation cell **1550**. Jobs agent **1520** may be communicatively coupled with an interface agent

1525 (typically a software module or application). Interface agent **1525** may include information processing and display software, such as a graphical user interface ("GUI"), an application program interface ("API"), or other interactive interface through which users and system processes can retrieve information about the status of storage operations. Through interface **1525**, users may optionally issue instructions to various storage operation cells **1550** regarding performance of the storage operations as described and contemplated by the present invention. For example, a user may modify a schedule concerning the number of pending snapshot copies or other types of copies scheduled as needed to suit particular needs or requirements. As another example, a user may employ the GUI to view the status of pending storage operations in some or all of the storage operation cells in a given network or to monitor the status of certain components in a particular storage operation cell (e.g., the amount of storage capacity left in a particular storage device).

Storage manager **1501** may also include a management agent **1530** that is typically implemented as a software module or application program. In general, management agent **1530** provides an interface that allows various management components **1501** in other storage operation cells **1550** to communicate with one another. For example, assume a certain network configuration includes multiple cells **1550** adjacent to one another or otherwise logically related in a WAN or LAN configuration (not shown). With this arrangement, each cell **1550** may be connected to the other through each respective interface agent **1525**. This allows each cell **1550** to send and receive certain pertinent information from other cells **1550** including status information, routing information, information regarding capacity and utilization, etc. These communication paths may also be used to convey information and instructions regarding storage operations.

For example, a management agent **1530** in first storage operation cell **1550** may communicate with a management agent **1530** in a second storage operation cell **1550** regarding the status of storage operations in the second storage operation cell. Another illustrative example includes the case where a management agent **1530** in first storage operation cell **1550** communicates with a management agent **1530** in a second storage operation cell to control the storage manager **1501** (and other components) of the second storage operation cell via the management agent **1530** contained in the storage manager **100**.

Another illustrative example is the case where management agent **130** in the first storage operation cell **1550** communicates directly with and controls the components in the second storage management cell **1550** and bypasses the storage manager **1501** in the second storage management cell. If desired, storage operation cells **1550** can also be organized hierarchically such that hierarchically superior cells control or pass information to hierarchically subordinate cells or vice versa.

Storage manager **1501** may also maintain an index, a database, or other data structure **1511**. The data stored in database **1511** may be used to indicate logical associations between components of the system, user preferences, management tasks, media containerization and data storage information or other useful data. For example, the storage manager **1501** may use data from database **1511** to track logical associations between media agent **1505** and storage devices **1515** (or movement of data as containerized from primary to secondary storage).

Generally speaking, a media agent, which may also be referred to as a secondary storage computing device, **1505** may be implemented as software module that conveys data, as

directed by storage manager **1501**, between a client computer **1585** and one or more storage devices **1515** such as a tape library, a magnetic media storage device, an optical media storage device, or any other suitable storage device. In one embodiment, secondary computing device **1505** may be communicatively coupled with and control a storage device **1515**. A secondary computing device **1505** may be considered to be associated with a particular storage device **1515** if that secondary computing device **1505** is capable of routing and storing data to particular storage device **1515**.

In operation, a secondary computing device **1505** associated with a particular storage device **1515** may instruct the storage device to use a robotic arm or other retrieval means to load or eject a certain storage media, and to subsequently archive, migrate, or restore data to or from that media. Secondary computing device **1505** may communicate with a storage device **1515** via a suitable communications path such as a SCSI or fiber channel communications link. In some embodiments, the storage device **1515** may be communicatively coupled to a data agent **105** via a Storage Area Network ("SAN").

Each secondary storage computing device **1505** may maintain a index, a database, or other data structure **1506** which may store index data generated during backup, migration, and restore and other storage operations as described herein. For example, performing storage operations on Microsoft Exchange data may generate index data. Such index data provides a secondary computing device **1505** or other external device with a fast and efficient mechanism for locating data stored or backed up. Thus, in some embodiments, a secondary storage computing device index **1506**, or a storage manager database **1511**, may store data associating a client **1585** with a particular secondary computing device **1505** or storage device **1515**, for example, as specified in a storage policy, while a database or other data structure in secondary computing device **1505** may indicate where specifically the client **1585** data is stored in storage device **1515**, what specific files were stored, and other information associated with storage of client **1585** data. In some embodiments, such index data may be stored along with the data backed up in a storage device **1515**, with an additional copy of the index data written to index cache in a secondary storage device. Thus the data is readily available for use in storage operations and other activities without having to be first retrieved from the storage device **1515**.

Generally speaking, information stored in cache is typically recent information that reflects certain particulars about operations that have recently occurred. After a certain period of time, this information is sent to secondary storage and tracked. This information may need to be retrieved and uploaded back into a cache or other memory in a secondary computing device before data can be retrieved from storage device **1515**. In some embodiments, the cached information may include information regarding format or containerization of archive or other files stored on storage device **1515**.

In some embodiments, certain components may reside and execute on the same computer. For example, in some embodiments, a client computer **1585** such as a data agent **1595**, or a storage manager **1501** coordinates and directs local archiving, migration, and retrieval application functions as further described in U.S. patent application Ser. No. 09/610,738. This client computer **1585** can function independently or together with other similar client computers **1585**.

Moreover, as shown in FIG. 15, clients **1585** and secondary computing devices **1505** may each have associated metabases (**1525** and **1560**, respectively). However in some embodiments each "tier" of storage, such as primary storage, sec-

ondary storage, tertiary storage, etc., may have multiple metabases or a centralized metabase, as described herein. For example, in FIG. 15, rather than a separate metabase **1525** associated with each client **1585**, the metabases on this storage tier may be centralized as discussed further herein. Similarly, second and other tiers of storage may have either centralized or distributed metabases. Moreover, mixed architectures systems may be used if desired, that may include a first tier centralized metabase system coupled to with a second tier storage system having distributed metabases and vice versa, etc.

Moreover, in operation, a storage manager **1501** or other management module may keep track of certain information that allows the storage manager to select, designated or otherwise identify metabases to be searched in response to certain queries as further described herein. Movement of data between primary and secondary storage may also involve movement of associated metadata and other tracking information as further described herein.

FIG. 16 is a diagram illustrating one arrangement of data that may be used in constructing an archive file according to one aspect of the invention. As shown, archive file **1600** may include header section **1610**, index section **1620** and payload section **1630**. Such an archive file may be constructed by a data agent at a client computing device when migrating data, for example, from primary to secondary storage, primary storage to other primary storage, etc. The payload section **1610** may include the data objects that are to be moved from a first location to a second location within the system (e.g., primary to secondary storage). These data objects may be identified by a data agent and designated to be moved pursuant to a storage preference such as a storage policy, a user preference, etc. Header **1610** may include routing and path information that identifies the origin and destination of the payload data and may include other information such as a list of files copied, checksums, etc. Index section **1620** may include certain other information regarding the payload data objects such as size, file type, and any offset or other logical indexing information that may be tracked by a storage management component or other component previously managing the data objects in the payload.

In some embodiments, storage managers may index information regarding archive files and related payload by time and storage on certain media so the archive files can be quickly located and/or retrieved. For example, it may be desired to identify certain data based on a query. The query may be analyzed and a certain time frame of interest may be identified. The system may use this information as a basis for a query search of certain index information (e.g., only search for records concerning operations that occurred during a specific time). This streamlines the search and retrieval process by narrowing the universe of data needs to be searched to locate responsive information.

FIG. 17 presents a flow chart **1700** that illustrates some of the steps that may be performed in moving data from primary storage to other storage devices within the system. First, at step **1702**, a query seeking certain data may be identified. The query may include aspects of data such as a schedule policy, storage policy, storage preference or other preference. The query may be analyzed and a primary metabase searched to identify data objects that satisfy the query (step **1704**). This may include parsing the query into constituent parts and analyzing each part alone or in combination with other portions as part of the evaluation process. At step **1706**, it may be determined, whether data objects satisfying the query are to be copied to other primary storage devices, to secondary storage devices or both (pursuant to a storage policy, etc.).

If at least some data objects satisfying the search criteria are to be copied to other primary storage devices, those data objects may be identified as further described herein and the target primary storage device(s) identified. This may involve consulting a storage policy or storage manager to determine the destination point. In some embodiments, destination may be determined dynamically, such that it is selected based on certain system preferences or optimization routines that select a storage device based on storage capacity, availability, data paths to the destination, etc.

At step **1708** the identified data objects may be copied from primary storage of a first computing device (the source) to primary storage of a second computing device (the target or destination). Any metadata associated with the first computing device describing the copied data may also be copied to a metabase associated with the second computing device such that this description information is not abandoned or lost, but rather travels with the copied data for subsequent use (step **1710**).

Next, at step **1712**, it may be determined whether the copied data objects and associated metadata are to be deleted from the source computing device. For example, this may be done in order to free storage space on the source computer or in accordance with certain data aging or migration criteria. If it is decided to delete the data objects (and associated metadata) the memory locations which include the data may be erased or designated for overwrite (step **1714** and **1716**).

In some embodiments the data objects may be deleted but certain metadata may be retained. If it is decided not delete the data objects, the data is retained and an index in an associated storage manager may be updated (step **1718**), for example by updating an index to reflect a new location, data object status, any changes, etc., and return to step **1702**. In other embodiments, if data is deleted from the system, for example, a user permanently deletes certain data from an application, that associated data may also be deleted from both primary and secondary storage devices and associated metabases to free storage space within the system.

Returning to step **1706**, it is also determined whether certain data objects currently stored in primary storage are to be migrated to one or more secondary storage devices. If so, an archive file similar to the one described in FIG. **16** or other data structure suitable for transport may be constructed or created by the source computing device with identified data objects placed in the payload section and header and index information added (step **1722**). Data may be moved from primary to secondary storage in predefined chunks which are constructed from such archive files, for example, using a data pipe, such as the data pipe described in U.S. Pat. No. 6,418,478 titled PIPELINED HIGH SPEED DATA TRANSFER MECHANISM.

Next, at step **1724** one or more target secondary storage devices may be identified. This may involve consulting a storage policy or storage manager to determine the destination point. In some embodiments, destination may be determined dynamically, such that it is selected based on certain system preferences or optimization routines that select a storage device based on storage capacity, availability, data paths to the destination, etc. Once the secondary storage device(s) are identified, the archive files may be routed to a media agent, storage manager, or other system component that supervises the transfer to the target secondary storage device (steps **1724** and **1728**). This may involve selecting and appropriate data transfer route and ensuring the proper resources and are available (e.g., bandwidth) such that the data may be copied with a certain period of time. Supervision may further include parsing a copy operation into several portions with each por-

tion being transferred by certain media agent or other resources, etc., to meet system or transfer requirements (e.g., a time window).

Next, the appropriate media within the target storage device may be identified (step **1730**) and the archive files may be transferred from the media management device to the secondary storage device (step **1732**). Such media may be selected from available media already associated with a similar data transfer or may be selected and reserved from an available media pool or scratch pool within the storage device. During or after the transfer, a media agent index or storage manager index associated with the secondary storage device may be updated to reflect the transfer (step **1733**). This may include copying the appropriate management files to the media management index such as offset, media ID file name or other management information.

At step **1734**, any metadata stored in a first metabase associated with the transferred data objects may also be transferred and used to update a second metabase associated with the target secondary storage device. Such metadata may be copied from the first metabase to the second metabase using network transmission resources. In some embodiments, the metadata in the first metabase may be deleted after it is confirmed the metadata has been copied to the second metabase. In other embodiments, the metadata may remain in both first and second metabases.

At step **1736**, it may be determined whether the data objects transferred from the primary storage device are to be deleted. If so, the data objects and associated metadata in a first metabase may be erased or otherwise designated for overwrite (steps **1738** and **1740**). In some cases, a pointer or other reference such as a file stub may be left in the original data location.

FIG. **18** presents a generalized block diagram of a hierarchically organized group of storage operation cells in a system to perform storage operations on electronic data in a computer network in accordance with an embodiment of the present invention. It will be understood that although the storage operation cells generally depicted in FIG. **18** have different reference numbers than the storage operation cell **1550** shown in FIG. **15**, these cells may be configured the same as or similar to the storage cell **1550** as depicted in FIG. **15**.

As shown, the system illustrated in FIG. **18** may include a master storage manager component **1835** and various other storage operations cells. As shown, the illustrative embodiment in FIG. **18** includes a first storage operation cell **1840**, a second storage operation cell **1845**, a third storage operation cell **1850**, a fourth storage operation cell **1855**, and may be extended to include nth storage operation cell, if desired (not shown). However, it will be understood this illustration is only exemplary and that fewer or more storage operation cells may be present or interconnected differently if desired.

Storage operation cells, such as the ones shown in FIG. **18** may be communicatively coupled and hierarchically organized. For example, a master storage manager **1835** may be associated with, communicate with, and direct storage operations for a first storage operation cell **1840**, a second storage operation cell **1845**, a third storage operation cell **1850**, and fourth storage operation cell **1855**. In some embodiments, the master storage manager **1835** may not be part of any particular storage operation cell. In other embodiments (not shown), master storage manager **1835** may itself be part of a certain storage operation cell. This logical organization provides a framework in which data objects, metadata and other management data may be hierarchically organized and associated with appropriate devices components (e.g., storage devices).

The storage operation cells may be configured in any suitable fashion, including those which involve distributed or centralized metabases. For example, storage operation cell **1840** may include a centralized primary storage metabase and a centralized secondary storage metabase, storage operation cell **1845** may include a centralized primary storage metabase and multiple secondary storage metabases, storage operation cell **1850** may include multiple primary storage metabases and a centralized secondary storage metabase, and storage operation cell **1855** may include multiple primary storage metabases and multiple secondary storage metabases (not shown). However, it will be understood that this is merely illustrative, and any other suitable configuration may be used if desired.

Thus, in operation, master storage manager **1835** may communicate with a management agent of the storage manager of the first storage operation cell **1840** (or directly with the other components of first cell **1840**) with respect to storage operations performed in the first storage operation cell **1840**. For example, in some embodiments, master storage manager **1835** may instruct the first storage operation cell **1840** with certain commands regarding a desired storage operation such as how and when to perform particular storage operations including the type of operation and the data on which to perform the operation.

Moreover, metabases associated with each storage operation cell may contain information relating to data and storage operations as described herein. In some embodiments, master storage manager **1835** may include a master metabase index or database (not shown) that reflects some or all of the meta-data information from the hierarchically subordinate storage operation cells within the system. This allows the system to consult the master storage index or database for information relating to data within those storage operation cells rather than requiring each cell be contacted or polled directly for such information.

In other embodiments, master storage manager **1835** may track the status of its associated storage operation cells, such as the status of jobs, system components, system resources, and other items, by communicating with manager agents (or other components) in the respective storage operation cells. Moreover, master storage manager **1835** may track the status of its associated storage operation cells by receiving periodic status updates from the manager agents (or other components) in the respective cells regarding jobs, system components, system resources, and other items. For example, master storage manager **1835** may use methods to monitor network resources such as mapping network pathways and topologies to, among other things, physically monitor storage operations and suggest, for example, alternate routes for storing data as further described herein.

In some embodiments, master storage manager **1835** may store status information and other information regarding its associated storage operation cells and other system information in an index cache, database or other data structure accessible to manager **1835**. A presentation interface included in certain embodiments of master storage manager **1835** may access this information and present it to users and system processes with information regarding the status of storage operations, storage operation cells, system components, and other information of the system.

In some embodiments, master storage manager **1835** may store and/or track metadata and other information regarding its associated storage operation cells and other system information in an index cache, database or other data structure accessible to manager **1835**. Thus, during a search procedure as further described herein, queries can be directed to a spe-

cific storage operation cell or cells based on the cell's function, past involvement, routing or other information maintained within the storage manager or other management component.

As mentioned above, storage operation cells may be organized hierarchically. With this configuration, storage operation cells may inherit properties from hierarchically superior storage operation cells or be controlled by other storage operation cells in the hierarchy (automatically or otherwise). Thus, in the embodiment shown in FIG. **18**, storage operation cell **1845** may control or is otherwise hierarchically superior to storage operation cells **1850** and **1855**. Similarly, storage operation cell **1850** may control storage operation cells **1855**. Alternatively, in some embodiments, storage operation cells may inherit or otherwise be associated with storage policies, storage preferences, storage metrics, or other properties or characteristics according to their relative position in a hierarchy of storage operation cells.

Storage operation cells may also be organized hierarchically according to function, geography, architectural considerations, or other factors useful or desirable in performing storage operations. For example, in one embodiment, storage operation cell **1840** may be directed to create snapshot copies of primary copy data, storage operation cell **1845** may be directed to create backup copies of primary copy data or other data. Storage operation cell **1840** may represent a geographic segment of an enterprise, such as a Chicago office, and storage operation cell **1845** may represent a different geographic segment, such as a New York office. In this example, the second storage operation cells **1845**, **1850** and **1855** may represent departments within the New York office. Alternatively, these storage operation cells could be further divided by function performing various types of copies for the New York office or load balancing storage operations for the New York office.

As another example, and as previously described herein, different storage operation cells directed to different functions may also contain the same or a subset of the same set of physical devices. Thus, one storage operation cell in accordance with the principles of the present invention may be configured to perform SRM operations and may contain the same, similar or a subset of the same physical devices as a cell configured to perform HSM or other types of storage operations. Each storage operation cell may, however, share the same parent or, alternatively, may be located on different branches of a storage operation cell hierarchy tree. For example, storage operation cell **1845** may be directed to SRM operations whereas storage operation cell **1855** may be directed to HSM operations. Those skilled in the art will recognize that a wide variety of such combinations and arrangements of storage operation cells are possible to address a broad range of different aspects of performing storage operations in a hierarchy of storage operation cells.

In some embodiments, hierarchical organization of storage operation cells facilitates, among other things, system security and other considerations. For example, in some embodiments, only authorized users may be allowed to access or control certain storage operation cells. For example, a network administrator for an enterprise may have access to many or all storage operation cells including master storage manager **1835**. But a network administrator for only the New York office, according to a previous example, may only have access to storage operation cells **1845-1855**, which form the New York office storage management system.

Moreover, queries performed by the system may be subject to similar restrictions. For example, depending on access privileges, users may be limited or otherwise excluded from

searching a certain cell or cells. For example, a user may be limited to searching information in cells or metabases within the system that are unrestricted or to those which specific access rights have been granted. For example, certain users may not have privileges to all information within the system. Accordingly, in some embodiments, as a default setting, users may have access privileges to information in cells that they interact with. Thus, confidential and sensitive information may be selectively restricted except only to certain users with express privileges (e.g., financial or legal information, etc.). For example, certain classification information within the metabases in the system may be restricted and therefore accessed only by those with the proper privileges.

Other restrictions on search criteria may include the scope of the search. For example, in a large network with many storage cells may require dedicating significant amounts of resources to perform go global or comprehensive searches. Thus, if a certain resource threshold is exceeded by a proposed search, the system may prompt that search to be modified or otherwise cancelled.

In other embodiments master storage manager **1835** may alert a user such as a system administrator when a particular resource is unavailable or congested. For example, a particular storage device might be full or require additional media. For example, a master storage manager may use information from an HSM storage operation cell and an SRM storage operation cell to present indicia or otherwise alert a user or otherwise identify aspects of storage associated with the storage management system and hierarchy of storage operation cells.

Alternatively, a storage manager in a particular storage operation cell may be unavailable due to hardware failure, software problems, or other reasons. In some embodiments, master storage manager **1835** (or another storage manager within the hierarchy of storage operation cells) may utilize the global data regarding its associated storage operation cells to suggest solutions to such problems when they occur (or act as a warning prior to occurrence). For example, master storage manager **1835** may alert the user that a storage device in a particular storage operation cell is full or otherwise congested, and then suggest, based on job and data storage information contained in its database, or associated metabase, or an alternate storage device. Other types of corrective actions based on such information may include suggesting an alternate data path to a particular storage device, or dividing data to be stored among various available storage devices as a load balancing measure or to otherwise optimize storage or retrieval time. In some embodiments, such suggestions or corrective actions may be performed automatically, if desired. This may include automatically monitoring the relative health or status of various storage operation cells and searching for information within the cells of the system relating to systems or resource performance within that cell (e.g., index, metabase, database, etc.) for use in diagnostics or for suggesting corrective action.

In alternate embodiments, HSM and SRM components may be aware of each other due to a common database or metabase of information that may include normalized data from a plurality of cells. Therefore, in those embodiments there is no need for such information to pass through a master storage manager as these components may be able to communicate directly with one another. For example, storage operation cell **1845** may communicate directly with storage operation cell **1855** and vice versa. This may be accomplished through a direct communications link between the two or by passing data through intermediate cells.

Moreover, in some embodiments searches may be performed across a numerous storage cells within the hierarchy. For example, a query may be posed to master storage manager **1835** that may pass the query down through the hierarchy from cells **1840** to **1845** to **1850** and **1855**, etc. This may be accomplished by passing the query from one manager component of each cell to another, or from one data classification agent to another, one metabase to another etc. The results may be passed upward through the hierarchy and compiled with other results such that master storage manager **1835** has a complete set of results to report. In other embodiments, each storage manager cell may report results directly to the requestor or to a designated location.

FIG. **19** presents a flow chart **1900** that illustrates some of the steps that may be involved in performing searches for data objects across systems that include multiple primary and secondary storage devices. First, at step **1905**, a query seeking certain data may be identified (e.g., from a storage policy, user preference, other process, etc.). The query may be analyzed to identify system components, such as clients potentially having information such as certain data objects or metadata that may satisfy the query (e.g., by excluding certain clients that are unlikely to have data being sought based on certain query parameters such as location, time frame, client or other component, department, application type, or any other criteria used to classify data as described herein, etc. (step **1910**)). Results may be presented based on a confidence factor indicating the likelihood that the results meet the specified parameters. For example, results substantially satisfying most or all criteria may be listed first with the confidence factors provided based on a percentage of the criteria satisfied (e.g., a query that returned results having three out of four criteria satisfied may be represented with a 75% confidence factor etc.). Less relevant results may be listed subsequently with the confidence factor provided based on any suitable relevant factor such as number of parameters satisfied, how close the match is, etc.

The search process may further involve consulting one or more indexes associated with the clients to identify where responsive data objects or other copies of client data, etc., may be located within the system. At step **1915**, it may be determined whether client data objects satisfying the query are located in primary storage, secondary storage, or both (e.g., based on index information in a storage manager). This may be based on polling various storage managers or a master storage manager that includes information the covers or represents whole system or the portion of system specified for search.

If it is determined that responsive data objects are only located on client(s) in primary storage, that client may be added to the list of clients to be searched (step **1955**). If it is determined that responsive data objects are located in secondary storage devices (or other primary storage locations that may be identified), the system may consult a storage manager index to identify archive files (or other files) based on certain query parameters such as a specified point in time, origination point, etc., or on index data stored in a storage manager index identifying archive files or other file associated with the data objects.

Next at step **1920**, storage managers may be consulted to identify responsive archive files. At step **1925**, media management components that may have handled responsive data objects are identified. This may be based on information retrieved from the storage manager index regarding archive files, e.g., an association of archive files with media agents and media items. It may then be determined whether the identified media management components have metadata

relating to the identified archive files available readily available in an index cache (step 1930).

This may be accomplished by searching for reference information relating to the identified archive files. If such information is already present in the cache, responsive data objects may be identified and retrieved using the index cache information, which may include, offsets and any file identifiers, etc., by the media management component, and the system may proceed to step 1940 (determine whether another media management component needs to be analyzed).

If not, the index information may need to be loaded from the secondary storage device so archive files may be retrieved and accessed. This may involve identifying the particular media on which the index data is stored and upload it to the media management component cache (step 1935). In some embodiments, a master storage manager or other component with information relating to files may be consulted to identify media containing the responsive information. These media may be mounted in drive or other input/output device and examined to locate the proper files or data structures. Index information may then be located and uploaded to an index or database associated with the appropriate media management component (e.g., media agent). This allows the media management component to locate and retrieve specific data objects on the media that satisfy the search criteria.

Next, if no further media management components have been identified, a list of media management components to be searched may be compiled (step 1945). At step 1950, a list of clients identified as potentially having responsive data objects may also be compiled. After a complete list of secondary storage devices and clients potentially having responsive data objects is identified, the associated metabascs are queried for these components, step 1960, and results are returned indicating data objects that may satisfy the search criteria, step 1965. In some embodiments, these results may be reviewed and analyzed to ensure relevance, with only reasonably relevant or responsive data objects actually being retrieved.

FIG. 20 presents a flow chart 2000 that illustrates some of the steps that may be involved in retrieving data objects from secondary storage (or other tiers or other storage locations) in accordance with principles of the present invention. This may be accomplished generally as follows. Certain data (e.g. data objects or associated metadata) from the system may need to be retrieved. That data may be requested and communicated to the system in the form of a query. The query may be used to search the system and identify media on which responsive data may be located. Once located, data satisfying the selection criteria may be uploaded and retrieved and analyzed for relevance, or other action may be taken. Or, alternatively, the identified data may be moved to other tiers of storage. More specific steps involved in this process may be as follows.

First, at step 2002, a query seeking certain data may be identified. The query may be analyzed to ascertain certain additional information that may assist in identifying responsive information such as identifying a certain point in time to search (step 2004). This may involve consulting storage manager and/or media agent index or database for responsive information relating to a certain point in time. This may also involve consulting certain metabascs for similar information that may be associated with these or other media management components providing copy and management functions. Point in time information may be specified by the user or may be assigned by the system absent a specific time frame established by the user. For example, a user may specify a certain time range within the query (e.g., a time range, a certain date, all information related to a project since its inception etc.).

The system however, may assign a certain time limit based on the query (e.g., such as based on the specifics of the query (e.g., only have data relating to a certain time frame)), and may limit the search to the time frame of information present in certain metabascs, master storage manager, or index within the system, and/or poll or otherwise communicate with storage devices within the system to determine the range or time frame of available data within the system and present the user with options for retrieving it (e.g., some, all within a time frame, etc.)

Next, at step 2006 certain archive files may be identified and associated media agents (step 2008) that may have been involved in transferring responsive data objects. This may be determined by consulting a master storage manager or other media management component index or metabase to determine whether the archive files have been handled by such components. Once the appropriate media agents have been identified, it may be determined whether information regarding the identified archive files is present in a cache or index associated with the media agents (step 2010). If not, the index information may need to be uploaded so the appropriate archive files may be retrieved and accessed. This process may be performed until all identified media agents have the appropriate index information loaded and/or until it is determined that no responsive information has been handled by the media agents and therefore no index information need be uploaded.

Next, at step 2016 data objects satisfying the query criteria may be identified by searching metabascs and/or indexes. In some embodiments, such data objects may be compiled into a list of data objects for present or subsequent retrieval. For example, such a list of responsive data objects may be provided to the user (which may itself satisfy the query) and then provide the user with the option to actually retrieve all or certain selected identified data objects.

At step 2018, the new destination for the data objects may be determined. For example, if certain data objects are being migrated off as part of an ILM operation, the query or other information may indicate the intent or reason for the search and the data object's destination. This may be useful in determining whether certain data objects are responsive to search criteria or query. At step 220 it may be determined whether the new destination is primary storage (a restore operation) or secondary or other tier of storage (ILM). Such information may be further useful in determining whether the data objects are likely to fall within a time frame or category of interest and thus may be useful in further identifying data objects of interest.

If the identified data objects are moving to other secondary storage tiers, the data objects may be repackaged into form suitable for secondary storage, which may include repackaging into an archive file, converting to a new format, compressing of the data objects and associated files, encryption, or any other containerization technique known in the art (step 2022).

Once the data objects are in a suitable format, they may be copied to the appropriate storage destination by the system. This may be accomplished by a media agent or media component in conjunction with a storage manager or other media management component that coordinate routing and the specifics involved with file transfer (step 2024), as further described herein. Metadata relating to the copied data objects may then be copied to a metabase associated with a computing device at the destination (step 2026).

For example, metadata relating to the data being copied may be copied along with the data to the secondary storage device and may be copied to an index in the media agent or other media management component involved in the data transfer. This allows the media management component to

locate and retrieve and otherwise manage the stored data. Such metadata may also be useful when performing searches of secondary storage devices (or other tiers) as further described herein. Metadata stored along with the data on the secondary storage device may be useful to restore or refresh the media agent index in the case of lost or corrupt data and also may be transferred along with the data on storage media in the case whether it is necessary to copy all such data (or actually physically relocate) to another storage device. A master storage manager index or metabase associated with destination computing device may be updated reflecting the arrival and new location of the transferred data objects and/or archive file for system management purposes (step 2034).

In some embodiments, the copied data objects and metadata may be deleted from the source location (steps 2028-2032). For example, at step 2028, it may be determined whether the copied data objects should be deleted based user preferences, storage policy requirements or other system constraints such as diminished storage capacity, etc. At steps 2030 and 2032 the data objects and records may be deleted. However, a stub, pointer or other referential element may be placed at the same logical location to act as a marker for the moved data. This allows subsequent operations to quickly track down and locate the moved data at its new location.

If, however, at step 2020, it is determined that the identified data objects are moving to primary storage, accordingly, the data objects may be reformatted (e.g., unpacked from archive file format) for copying to a computing device (step 2038). Next the unpacked data may be copied to a target computing device along with any associated metadata (steps 2040 and 2042). For example, this may involve reading metadata and/or index information from the archive file and repopulating the metabase and/or management component indexes with this information as further described herein. For example, metadata from the archive file may be retrieved and integrated into a metabase associated with the target computing device including information relating to data management and as well as certain content and storage information as further described herein with respect to the classification process and metabase population. Thus, such archive information may be fully restored to primary storage and any associated information, such as metabase information may be searched and retrieved accordingly.

Moreover, information relating to system management may be uploaded and used to repopulate storage management components within the system such as a storage manager or master storage manager reflecting the return of the retrieved data to primary storage (step 2050). For example, a storage manager index may be updated to reflect the presence of the retrieved data along with certain management information such as logical offsets and location of the retrieved information such that the retrieved information may be located and accessed. Other management components, such as a master storage manager may also be updated with the appropriate identification and location information to reflect the return of the retrieved data within the system.

In certain embodiments, the copied data and metadata may be deleted from the source location (steps 2044-2048). For example, at step 2044, it may be determined whether the copied data objects in secondary storage should be deleted based user preferences, storage policy requirements or other system constraints such as diminished storage capacity, etc. At steps 2046 and 2048 the data objects and records may be deleted within the system including any metabase or other system management information associated with the retrieved data. Storage management components such as stor-

age managers, media agents may also be updated to reflect the removal or deletion of such information (step 2050).

Systems and modules described herein may comprise software, firmware, hardware, or any combination(s) of software, firmware, or hardware suitable for the purposes described herein. Software and other modules may reside on servers, workstations, personal computers, computerized tablets, personal digital assistants (PDAs), and other devices suitable for the purposes described herein. In other words, the software and other modules described herein may be executed by a general-purpose computer, e.g., a server computer, wireless device or personal computer. Those skilled in the relevant art will appreciate that aspects of the invention can be practiced with other communications, data processing, or computer system configurations, including: Internet appliances, handheld devices (including PDAs), wearable computers, all manner of cellular or mobile phones, multi-processor systems, microprocessor-based or programmable consumer electronics, set-top boxes, network PCs, mini-computers, mainframe computers, and the like. Indeed, the terms "computer," "server," "host," "host system," and the like are generally used interchangeably herein, and refer to any of the above devices and systems, as well as any data processor. Furthermore, aspects of the invention can be embodied in a special purpose computer or data processor that is specifically programmed, configured, or constructed to perform one or more of the computer-executable instructions explained in detail herein.

Software and other modules may be accessible via local memory, via a network, via a browser or other application in an ASP context, or via other means suitable for the purposes described herein. Aspects of the invention can also be practiced in distributed computing environments where tasks or modules are performed by remote processing devices, which are linked through a communications network, such as a Local Area Network (LAN), Wide Area Network (WAN), or the Internet. In a distributed computing environment, program modules may be located in both local and remote memory storage devices. Data structures described herein may comprise computer files, variables, programming arrays, programming structures, or any electronic information storage schemes or methods, or any combinations thereof, suitable for the purposes described herein. User interface elements described herein may comprise elements from graphical user interfaces, command line interfaces, and other interfaces suitable for the purposes described herein. Screenshots presented and described herein can be displayed differently as known in the art to input, access, change, manipulate, modify, alter, and work with information.

Aspects of the invention may be stored or distributed on computer-readable media, including magnetically or optically readable computer discs, hard-wired or preprogrammed chips (e.g., EEPROM semiconductor chips), nanotechnology memory, biological memory, or other data storage media. Indeed, computer implemented instructions, data structures, screen displays, and other data under aspects of the invention may be distributed over the Internet or over other networks (including wireless networks), on a propagated signal on a propagation medium (e.g., an electromagnetic wave(s), a sound wave, etc.) over a period of time, or they may be provided on any analog or digital network (packet switched, circuit switched, or other scheme).

Unless the context clearly requires otherwise, throughout the description and the claims, the words "comprise," "comprising," and the like are to be construed in an inclusive sense, as opposed to an exclusive or exhaustive sense; that is to say, in the sense of "including, but not limited to." As used herein, the terms "connected," "coupled," or any variant thereof,

means any connection or coupling, either direct or indirect, between two or more elements; the coupling of connection between the elements can be physical, logical, or a combination thereof. Additionally, the words "herein," "above," "below," and words of similar import, when used in this application, shall refer to this application as a whole and not to any particular portions of this application. Where the context permits, words in the above Detailed Description using the singular or plural number may also include the plural or singular number respectively. The word "or," in reference to a list of two or more items, covers all of the following interpretations of the word: any of the items in the list, all of the items in the list, and any combination of the items in the list.

The above detailed description of embodiments of the invention is not intended to be exhaustive or to limit the invention to the precise form disclosed above. While specific embodiments of, and examples for, the invention are described above for illustrative purposes, various equivalent modifications are possible within the scope of the invention, as those skilled in the relevant art will recognize. For example, while processes or blocks are presented in a given order, alternative embodiments may perform routines having steps, or employ systems having blocks, in a different order, and some processes or blocks may be deleted, moved, added, subdivided, combined, and/or modified to provide alternative or subcombinations. Each of these processes or blocks may be implemented in a variety of different ways. Also, while processes or blocks are at times shown as being performed in series, these processes or blocks may instead be performed in parallel, or may be performed at different times.

The teachings of the invention provided herein can be applied to other systems, not necessarily the system described above. The elements and acts of the various embodiments described above can be combined to provide further embodiments. Any patents and applications and other references noted above, including any that may be listed in accompanying filing papers, are incorporated herein by reference. Aspects of the invention can be modified, if necessary, to employ the systems, functions, and concepts of the various references described above to provide yet further embodiments of the invention.

We claim:

1. A system for distributing metadata related to data stored in a network, comprising:

- a first storage device containing first data accessible via the network;
- a second storage device containing second data accessible via the network;
- a first metabase containing metadata and having a first characteristic, wherein the metadata includes classifications of the data stored on the first and second storage devices, the classifications including information used to facilitate performance of a data storage operation on the first data or on the second data; and
- a second metabase containing metadata and having a second characteristic, wherein the metadata includes classifications of the data stored on the first and second storage devices, the classifications including information used to facilitate performance of a data storage operation on first data or on the second data; and

a metabase manager component in communication with the first metabase and the second metabase over the network, wherein the metabase manager component is programmed to direct requests to perform data storage operations to the first metabase or the second metabase based on the characteristics of the metabases, wherein the metabase manager component directs a request to perform a data storage operation by:

- determining the first characteristic of the first metabase and the second characteristic of the second metabase;
- evaluating the first and second characteristics against one or more criteria; and

directing the request to perform the data storage operation to the first metabase or the second metabase based on the evaluation.

2. The system of claim 1 wherein the first metabase contains metadata classifying the data stored on the first storage device and the second metabase contains metadata classifying the data stored on the second storage device.

3. The system of claim 1 wherein both the first and second metabases contain copies of metadata classifying the data stored on the first and second storage devices.

4. The system of claim 1, wherein the metabase manager component directs requests to perform data storage operations based at least in part on a past level of activity of the metabases.

5. The system of claim 1 wherein the metabase manager component directs requests to perform data storage operations based at least in part on a current level of activity of the metabases.

6. The system of claim 1 wherein the metabase manager component directs requests to perform data storage operations based at least in part on a projected level of activity of the metabases.

7. The system of claim 1 wherein the metabase manager component directs requests to perform data storage operations based on the data to be accessed by the storage request.

8. The system of claim 1 wherein the metabase manager component directs requests to perform data storage operations based on access control information for the metabases.

9. The system of claim 1 wherein the first and second characteristics of the metabases are associated with a location of the metabases.

10. The system of claim 1 wherein the first and second characteristics of the metabases are associated with a hierarchy of departments within a business organization that stores data to the first and second storage devices.

11. The system of claim 1 wherein the first and second characteristics of the metabases are associated with subsidiaries of a business organization that stores data to the first and second storage devices.

12. The method of claim 1 wherein the metabase manager component determines the first and second characteristics by monitoring the first and second metabases.

13. The method of claim 1 wherein the metabase manager component evaluates the first and second characteristics against one or more criteria relating to availability of the first and second metabases.

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 7,660,800 B2
APPLICATION NO. : 11/564153
DATED : February 9, 2010
INVENTOR(S) : Anand Prahlad et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In column 1, line 45, after "10/818,749," insert -- titled --.

In column 1, line 49, after "10/877,831," insert -- titled --.

In column 2, line 6, delete "2005," and insert -- 2005. --, therefor.

In column 11, line 12, after "using" delete "a".

In column 21, line 7-8, delete "classifications" and insert -- classification --, therefor.

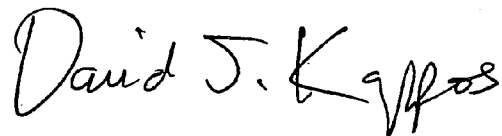
In column 21, line 14, after "and" insert -- / --.

In column 27, line 66, after "However" insert -- , --.

In column 30, line 33, after "location" insert -- . --.

Signed and Sealed this

Twenty-ninth Day of June, 2010

A handwritten signature in black ink that reads "David J. Kappos". The signature is written in a cursive, flowing style with a large initial "D" and a stylized "K".

David J. Kappos
Director of the United States Patent and Trademark Office

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 7,660,800 B2
APPLICATION NO. : 11/564153
DATED : February 9, 2010
INVENTOR(S) : Anand Prahlad et al.

Page 1 of 1

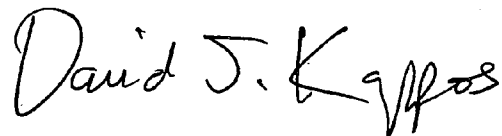
It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

On the title page, Item (*) Notice should read:

Subject to any disclaimer, the term of this patent is extended or adjusted under 35
U.S.C. 154(b) by 217 days.

Signed and Sealed this

Thirty-first Day of August, 2010

A handwritten signature in black ink that reads "David J. Kappos". The signature is written in a cursive, flowing style with a large, stylized 'D' and 'K'.

David J. Kappos
Director of the United States Patent and Trademark Office